



Red Hat JBoss A-MQ 6.1

Fabric Guide

A system for provisioning containers deployed across a network

Red Hat JBoss A-MQ 6.1 Fabric Guide

A system for provisioning containers deployed across a network

JBoss A-MQ Docs Team

Content Services

fuse-docs-support@redhat.com

Legal Notice

Copyright © 2014 Red Hat.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution-Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

Fabric enables you to install, start up, and provision remote containers across a network with support for centralized, highly available container configuration, based on Apache Zookeeper.

Table of Contents

CHAPTER 1. AN OVERVIEW OF FUSE FABRIC	4
1.1. CONCEPTS	4
1.2. CONTAINERS	6
1.3. PROVISIONING AND CONFIGURATION	7
PART I. BASIC FABRIC DEPLOYMENT	8
CHAPTER 2. GETTING STARTED WITH FUSE FABRIC	9
2.1. CREATE A FABRIC	9
2.2. DEPLOY A PROFILE	10
2.3. UPDATE A PROFILE	11
2.4. SHUTTING DOWN THE CONTAINERS	12
CHAPTER 3. CREATING A NEW FABRIC	14
STATIC IP ADDRESS REQUIRED FOR FABRIC SERVER	14
PROCEDURE	14
FABRIC CREATION PROCESS	16
EXPANDING A FABRIC	16
CHAPTER 4. FABRIC CONTAINERS	18
4.1. CHILD CONTAINERS	18
4.2. SSH CONTAINERS	19
4.3. FABRIC CONTAINERS ON WINDOWS	21
4.4. CLOUD CONTAINERS	23
CHAPTER 5. FABRIC PROFILES	29
5.1. INTRODUCTION TO PROFILES	29
5.2. WORKING WITH PROFILES	30
5.3. PROFILE VERSIONS	36
CHAPTER 6. FABRIC8 MAVEN PLUG-IN	38
6.1. PREPARING TO USE THE PLUG-IN	38
6.2. USING THE PLUG-IN TO DEPLOY A MAVEN PROJECT	38
6.3. CONFIGURING THE PLUG-IN	40
6.4. CONFIGURATION PROPERTIES	42
CHAPTER 7. ACTIVEMQ BROKERS AND CLUSTERS	44
7.1. CREATING A STANDALONE BROKER INSTANCE	44
7.2. CONNECTING TO A BROKER	46
7.3. TOPOLOGIES	47
7.4. BROKER CONFIGURATION	53
PART II. FABRIC IN PRODUCTION	62
CHAPTER 8. FABRIC ENSEMBLE AND REGISTRY	63
8.1. FABRIC REGISTRY	63
8.2. ADMINISTERING A FABRIC ENSEMBLE	64
CHAPTER 9. FABRIC AGENTS	66
9.1. INTRODUCTION	66
9.2. THE CONFIGURATION ADMIN BRIDGE	66
9.3. THE DEPLOYMENT AGENT	67
CHAPTER 10. ALLOCATING PORTS	70
10.1. THE PORT SERVICE	70

10.2. USING THE PORT SERVICE	72
CHAPTER 11. GATEWAY	76
11.1. GATEWAY ARCHITECTURE	76
11.2. RUNNING THE GATEWAY	76
11.3. CONFIGURING THE GATEWAY	77
11.4. VERSIONING	78
11.5. URI TEMPLATE EXPRESSIONS	79
CHAPTER 12. SECURING FABRIC CONTAINERS	81
DEFAULT AUTHENTICATION SYSTEM	81
MANAGING USERS	81
OBFUSCATING STORED PASSWORDS	81
ENABLING LDAP AUTHENTICATION	82
CHAPTER 13. CONFIGURING A FABRIC'S MAVEN PROXY	83
OVERVIEW	83
DEFAULT REPOSITORIES	83
CHANGING THE REPOSITORIES	83
CHAPTER 14. OFFLINE REPOSITORIES	85
14.1. OFFLINE REPOSITORY FOR A PROFILE	85
14.2. OFFLINE REPOSITORY FOR A VERSION	85
14.3. OFFLINE REPOSITORY FOR A MAVEN PROJECT	85
CHAPTER 15. PATCHING	87
15.1. PATCHING A CONTAINER IN A FABRIC	87
APPENDIX A. EDITING PROFILES WITH THE BUILT-IN TEXT EDITOR	90
A.1. EDITING AGENT PROPERTIES	90
A.2. EDITING OSGI CONFIG ADMIN PROPERTIES	93
A.3. EDITING OTHER RESOURCES	94
A.4. PROFILE ATTRIBUTES	96
APPENDIX B. FABRIC URL HANDLERS	98
B.1. PROFILE URL HANDLER	98
B.2. ZK URL HANDLER	98
B.3. BLUEPRINT URL HANDLER	98
B.4. SPRING URL HANDLER	99
APPENDIX C. PROFILE PROPERTY RESOLVERS	100
C.1. SUBSTITUTING SYSTEM PROPERTIES	100
C.2. SUBSTITUTING ENVIRONMENT VARIABLES	100
C.3. SUBSTITUTING CONTAINER ATTRIBUTES	101
C.4. SUBSTITUTING PID PROPERTIES	102
C.5. SUBSTITUTING ZOOKEEPER NODE CONTENTS	103
C.6. CHECKSUM PROPERTY RESOLVER	104
C.7. PORT PROPERTY RESOLVER	104
APPENDIX D. TECHNOLOGY-SPECIFIC DISCOVERY MECHANISMS	105
D.1. ACTIVEMQ ENDPOINT DISCOVERY	105
D.2. CAMEL ENDPOINT DISCOVERY	105
D.3. CXF ENDPOINT DISCOVERY	106
D.4. OSGI SERVICE DISCOVERY	106

CHAPTER 1. AN OVERVIEW OF FUSE FABRIC

Abstract

Fuse Fabric is a lightweight runtime environment that focuses on centralizing provisioning and configuration spanning from small environments, limited to a few JVMs or systems, to larger production environments which utilize an open hybrid cloud consisting of containers on both cloud services and physical hosts. The Fuse Fabric technology layer supports the scalable deployment of JBoss Fuse containers across a network, and it enables a variety of advanced features, such as remote installation and provisioning of containers, phased rollout of new versions of libraries and applications, load-balancing, and failover of deployed endpoints.

1.1. CONCEPTS


Fabric

A *fabric* is a collection of containers that share a *fabric registry*, where the fabric registry is a replicated database that stores all information related to provisioning and managing the containers. A fabric is intended to manage a distributed network of containers, where the containers are deployed across multiple hosts.

A sample fabric

Figure 1.1, “Containers in a Fabric” shows an example of a distributed collection of containers that belong to a single fabric.

Figure 1.1. Containers in a Fabric



Registry

JBoss Fuse Fabric uses [Apache ZooKeeper](#) (a highly reliable, distributed coordination service) as its registry to store the cluster configuration and the node registration data.

ZooKeeper is designed to provide data consistency and high availability across data centers. ZooKeeper protects against network splits using a quorum of ZooKeeper servers: for example, you might run five ZooKeeper servers and so long as you have quorum (that is, at least three out of the five servers available) you are reliable and are not affected by a network split.

Conceptually, the Fabric registry consists of two main parts:

- *Configuration Registry*—the logical configuration of your fabric, which typically contains no physical machine information. It contains details of the applications to be deployed and their dependencies.
- *Runtime Registry*—contains details of how many machines are actually running, their physical location, and what services they are implementing.

Fabric Ensemble

A *Fabric Ensemble* is a collection of Fabric Servers that collectively maintain the state of the fabric registry. The Fabric Ensemble implements a replicated database and uses a [quorum-based voting system](#) to ensure that data in the fabric registry remains consistent across all of the fabric's containers. To guard against network splits in a quorum-based system, it is a requirement that *the number of Fabric Servers in a Fabric Ensemble is always an odd number*.

The number of Fabric Servers in a fabric is typically 1, 3, or 5. A fabric with just one Fabric Server is suitable for experimentation only. A live production system should have at least 3 or 5 Fabric Servers, installed on separate hosts, to provide fault tolerance.

Fabric Server

A Fabric Server has a special status in the fabric, because it is responsible for maintaining a replica of the fabric registry. In each Fabric Server, a registry service is installed (labeled R in [Figure 1.1, “Containers in a Fabric”](#)). The registry service (based on Apache ZooKeeper) maintains a replica of the registry database and provides a ZooKeeper server, which ordinary agents can connect to in order to retrieve registry data.

Fabric Container (managed container)

A *Fabric container* (or *managed container*) is aware of the locations of all of the Fabric Servers, and it can retrieve registry data from any Fabric Server in the Fabric Ensemble. A *Fabric agent* (labeled A in [Figure 1.1, “Containers in a Fabric”](#)) is installed in each Fabric container. The Fabric Agent actively monitors the fabric registry, and whenever a relevant modification is made to the registry, it immediately updates its container to keep the container consistent with the registry settings.

Fabric Agent

Fabric defines a provisioning agent or *Fabric agent*, which relies on *profiles*. The Fabric agent runs on each managed container and its role is to provision the container according to the profiles assigned to it. The Fabric agent retrieves the configuration, bundles and features (as defined in the profile overlay), calculates what needs to be installed (or uninstalled) and, finally, performs the required actions.

The Fabric agent does not just provision applications; it is also capable of provisioning Fabric and the OSGi framework.

The Fabric agent retrieves any required Maven artifacts from the Maven repositories specified by its profile, which are accessed through the Maven proxies managed by the fabric.

Git

JBoss Fuse Fabric has *git* as the distributed version control mechanism for all configurations allowing for a full audit history; in addition, all changes are versioned and replicated onto each machine. By leveraging a well known technology users can easily perform diffs, merges, and continuous integration.

Profile

A Fabric *profile* is an abstract unit of deployment, capable of holding all of the data required for deploying an application into a Fabric Container. Profiles are used exclusively in the context of fabrics.


A profile consists of a collection of OSGi bundles and Karaf features to be provisioned, and a list of configurations for the OSGi Configuration Administration service. Multiple profiles can be associated with a given container, allowing the container to serve multiple purposes.

In theory, if you have a fabric where each managed container runs exactly the same set of features, with exactly the same configuration, you need only *one* profile. However, in most real-world use cases, you need to provision different features across the nodes, or at least have different configurations.

For example, you might want to run different kinds of application on different nodes: message brokers, web applications, ESBs, transformation engines, or proxies. Each kind of application can be defined by a profile that you manage as a single logical entity, irrespective of how many instances are deployed. Alternatively, you might want to run the same features everywhere, but with location specific configurations. You might deploy message brokers in different locations, or cache services in different data centres or geographical regions.

Profiles can also use inheritance, so that parts of a configuration can be shared across multiple profiles (which is conceptually similar to the use of trees in an LDAP repository). The aggregate of configuration settings is computed using an overlay mechanism, which allows a profile to override values from its parents. This approach provides power and flexibility, while avoiding the unnecessary repetition of configuration values.

Profiles are stored in ZooKeeper, which means that they are automatically and immediately propagated to all nodes in the fabric.



IMPORTANT

The presence of a Fabric agent in a container completely changes the deployment model, *requiring you to use profiles exclusively* as the unit of deployment. Although it is still possible to deploy an individual bundle or feature (using `osgi:install` or `features:install`, respectively), these modifications are impermanent. As soon as you restart the container or refresh its contents, the Fabric agent replaces the container's existing contents with whatever is specified by the deployed profiles.

1.2. CONTAINERS

Fuse Fabric is designed to have a number of containers, each of which may be running on a different platform. For instance, all of the following are supported:

- Java processes running directly on hardware

- Java processes running directly on hardware.
- A PaaS (Platform as a Service) such as OpenShift for either a public or private cloud.
- OpenStack as an IaaS (Infrastructure as a Service).
- Amazon Web Services, Rackspace or some other IaaS to manage services.
- Docker containers for each service.
- An open hybrid cloud composing all of the above entries.

1.3. PROVISIONING AND CONFIGURATION

Overview

Each node in the cluster has a logical name and on startup the provisioning agent registers an ephemeral node in ZooKeeper. It then looks in the configuration using its logical name for its profile of what configuration to use and what services to provision. If there is no configuration or the logical name it uses the default profile.

Each node watches the relevant parts of the ZooKeeper tree for changes; so that as soon as any profile is updated the dependent nodes refresh and re-apply any provisioning or configuration changes.

Changing the configuration

You can use the Fabric command line shell to modify a profile's configuration at run time, which means you can dynamically control configuration and provisioning.

Profiles make it possible to control a group of nodes (or all nodes) in a single operation. For example, you could set global configuration properties that affect all nodes, or properties that affect only a group of nodes (for example, affecting only message brokers, or affecting only web servers).

How discovery works

Each container registers details of its machine, host name and connection URLs (for example, its JMX URL) with the ZooKeeper runtime registry. Consequently, it is relatively easy to discover a node's location and its connection details by introspecting the fabric. For example, to perform operations on a remote fabric container from within the fuse console, all that you need to know is the container's name.

In addition, Fabric supports a number of discovery mechanisms for discovering application endpoints and services in the fabric—see [Appendix D, *Technology-specific Discovery Mechanisms*](#) for details.

PART I. BASIC FABRIC DEPLOYMENT

Abstract

Get started with Fuse Fabric and learn how to perform basic administration tasks.

CHAPTER 2. GETTING STARTED WITH FUSE FABRIC

Abstract

This tutorial provides basic information and explains how to set up the simplest Fabric system, by creating some containers that run on your local machine and then deploying an example profile to a child container.


Additional information on setting up a Fabric is covered in more detail in both [Chapter 3, *Creating a New Fabric*](#) and [Section 4.1, “Child Containers”](#).

2.1. CREATE A FABRIC

Overview

[Figure 2.1](#) shows an overview of a sample fabric that you will create. The Fabric Ensemble consists of just one Fabric Server (making this fabric suitable only for experimental use) and two managed child containers.

Figure 2.1. A Sample Fabric with Child Containers



Steps to create the fabric

To create the simple fabric shown in [Figure 2.1, “A Sample Fabric with Child Containers”](#), follow these steps:

1. To create the first fabric container, which acts as the seed for the new fabric, enter this console command:

```
JBossFuse:karaf@root> fabric:create --new-user AdminUser --new-user-
password AdminPass
```

```
--zookeeper-password ZooPass
--resolver manualip --manual-ip 127.0.0.1 --wait-for-provisioning
```

The current container, named **root** by default, becomes a Fabric Server with a registry service installed. Initially, this is the only container in the fabric. The `--new-user` and `--new-user-password` options specify the credentials for a new administrator user. The Zookeeper password is used to protect sensitive data in the Fabric registry service (all of the nodes under `/fabric`). The `--manual-ip` option specifies the loopback address, `127.0.0.1`, as the Fabric Server's IP address.



NOTE

A Fabric Server requires a static IP address. For simple trials and tests, you can use the loopback address, `127.0.0.1`, to work around this requirement. But if you are deploying a fabric in production or if you want to create a distributed ensemble, you *must* assign a static IP address to each of the Fabric Server hosts.



NOTE

Most of the time, you are not prompted to enter the Zookeeper password when accessing the registry service, because it is cached in the current session. When you *join a container to a fabric*, however, you must provide the fabric's Zookeeper password.

2. Create a child container. Assuming that your root container is named **root**, enter this console command:

```
JBossFuse:karaf@root> fabric:container-create-child root child
The following containers have been created successfully:
Container: child.
```

3. Invoke the following command to monitor the status of the child container, as it is being provisioned:

```
JBossFuse:karaf@root> shell:watch container-list
```

After the deployment of the **child** has completed, you should see a listing something like this:

```
JBossFuse:karaf@root> shell:watch container-list
[id]                                [version] [alive] [profiles]
[provision status]
root                                1.0        true   fabric, fabric-
ensemble-0000-1, fuse-esb-full success
  child                               1.0        true   default
success
```

Type the **Return** key to get back to the JBoss Fuse console prompt.

2.2. DEPLOY A PROFILE

Deploy a profile to the child container

Having created the child container, as described in [Section 2.1, “Create a Fabric”](#), you can now deploy the a profile to it. To do so, follow these steps:

1. Deploy the `example-camel-hello` profile into the `child` container by entering this console command:

```
JBossFuse:karaf@root> fabric:container-change-profile child example-camel-hello
```

2. Verify that the `example-camel-hello` profile deploys successfully to the `child` container, using the `fabric:container-list` command. Enter the following command to monitor the container status:

```
JBossFuse:karaf@root> shell:watch container-list
```

And wait until the `child` container status changes to `success`.

View the sample output

When it is running, the `example-camel-hello` profile writes a message to the container's log every five seconds. To verify that the profile is running properly, you can look for these messages in the child container's log, as follows:

1. Connect to the `child` container, by entering the following console command:

```
JBossFuse:karaf@root> container-connect child
```

2. After logging on to the `child` container, view the `child` container's log using the `log:tail` command, as follows:

```
JBossFuse:karaf@root> log:tail
```

You should see some output like the following:

```
■
```

3. Type Ctrl-C to exit the log view and get back to the child container's console prompt.

2.3. UPDATE A PROFILE

Atomic container upgrades

Normally, when you edit a profile that is already deployed in a container, *the modification takes effect immediately*. This is because the Fabric Agent in the affected container (or containers) actively monitors the fabric registry in real time.

In practice, however, immediate propagation of profile modifications is often undesirable. In a production system, you typically want to roll out changes incrementally: for example, initially trying out the change on just one container to check for problems, before you make changes globally to all containers. Moreover, sometimes several edits must be made together to reconfigure an application in a consistent way.

Profile versioning

For quality assurance and consistency, it is typically best to modify profiles *atomically*, where several modifications are applied simultaneously. To support atomic updates, fabric implements profile versioning. Initially, the container points at version 1.0 of a profile. When you create a new profile version (for example, version 1.1), the changes are invisible to the container until you upgrade it. After you are finished editing the new profile, you can apply all of the modifications simultaneously by upgrading the container to use the new version 1.1 of the profile.

Upgrade to a new profile

For example, to modify the `example-camel-hello` profile, when it is deployed and running in a container, follow the recommended procedure:

1. Create a new version, 1.1, to hold the pending changes by entering this console command:

```
JBossFuse:karaf@root> fabric:version-create
Created version: 1.1 as copy of: 1.0
```

The new version is initialised with a copy of all of the profiles from version 1.0.

2. Use the `fabric:profile-edit` command to change the message that is written to the container log by the Camel route. Enter the following `profile-edit` command to edit the `camel.xml` resource:

```
JBossFuse:karaf@root> fabric:profile-edit --resource camel.xml
example-camel-hello 1.1
```

This opens the built-in text editor for editing profile resources (see [Appendix A, Editing Profiles with the Built-In Text Editor](#)).

Remember to specify version `1.1` to the `fabric:profile-edit` command, so that the modifications are applied to version 1.1 of the `example-camel-hello` profile.

3. Upgrade the `child` container to version `1.1` by entering this console command:

```
JBossFuse:karaf@root> fabric:container-upgrade 1.1 child
```

Roll back to an old profile

You can easily roll back to the old version of the `example-camel-hello` profile, using the `fabric:container-rollback` command like this:

```
JBossFuse:karaf@root> fabric:container-rollback 1.0 child
```

2.4. SHUTTING DOWN THE CONTAINERS

Shutting down the containers

Because the child containers run in their own JVMs, they do *not* automatically stop when you shut down the root container. To shut down a container and its children, first stop its children using the `fabric:container-stop` command. For example, to shut down the current fabric completely, enter

these console commands:

```
JBossFuse:karaf@root> fabric:container-stop child  
JBossFuse:karaf@root> shutdown -f
```

After you restart the root container, you must explicitly restart the children using the **fabric:container-start** console command.

CHAPTER 3. CREATING A NEW FABRIC

Abstract

When there are no existing fabric's to join, or you want to start a new fabric, you can create a new one from a standalone container.

STATIC IP ADDRESS REQUIRED FOR FABRIC SERVER

The IP address and hostname associated with the Fabric Servers in the Fabric ensemble are of critical importance to the fabric. Because these IP addresses and hostnames are used for configuration and service discovery (through the Zookeeper registry), they *must not change* during the lifetime of the fabric.

You can take either of the following approaches to specifying the IP address:

- For simple examples and tests (with a single Fabric Server) you can work around the static IP requirement by using the loopback address, `127.0.0.1`.
- For distributed tests (multiple Fabric Servers) and production deployments, you *must* assign a static IP address to each of the Fabric Server hosts.



WARNING

Beware of volatile IP addresses resulting from VPN connections, WiFi connections, and even LAN connections. If a Fabric Server binds to one of these volatile IP addresses, it will cease to function after the IP address has gone away. It is recommended that you always use the `--resolver manualip --manual-ip StaticIPAddress` options to specify the static IP address explicitly, when creating a new Fabric Server.

PROCEDURE

To create a new fabric:

1. (Optional) Customise the name of the root container by editing the `InstallDir/etc/system.properties` file and specifying a different name for this property:

```
karaf.name=root
```



NOTE

For the first container in your fabric, this step is optional. But at some later stage, if you want to join a root container to the fabric, you might need to customise the container's name to prevent it from clashing with any existing root containers in the fabric.

- Any existing users in the `InstallDir/etc/users.properties` file are automatically used to initialize the fabric's user data, when you create the fabric. You can populate the `users.properties` file, by adding one or more lines of the following form:

```
Username=Password[,RoleA][,RoleB]...
```

But there must *not* be any users in this file that have the `admin` role. If the `InstallDir/etc/users.properties` already contains users with the `admin` role, you should *delete those users* before creating the fabric.



WARNING

If you leave some administrator credentials in the `users.properties` file, this represents a security risk because the file could potentially be accessed by other containers in the fabric.



NOTE

The initialization of user data from `users.properties` happens only once, at the time the fabric is created. After the fabric has been created, any changes you make to `users.properties` will have *no effect* on the fabric's user data.

- If you use a VPN (virtual private network) on your local machine, it is advisable to log off VPN *before* you create the fabric and to *stay logged off* while you are using the local container.



NOTE

A local Fabric Server is permanently associated with a fixed IP address or hostname. If VPN is enabled when you create the fabric, the underlying Java runtime is liable to detect and use the VPN hostname instead of your permanent local hostname. This can also be an issue with multi-homed machines.

- Start up your local container.

In JBoss A-MQ, start the local container as follows:

```
cd InstallDir/bin
./amq
```



NOTE

If you want to create a fabric on a remote container, you can log into it using the `InstallDir/bin/client` command-line utility.

- Create a new fabric by entering the following command:

```
JBossFuse:karaf@root> fabric:create --new-user AdminUser --new-user-
```

```
password AdminPass
--zookeeper-password ZooPass
--resolver manualip --manual-ip StaticIPAddress --wait-for-
provisioning
```

The current container, named `root` by default, becomes a Fabric Server with a registry service installed. Initially, this is the only container in the fabric. The `--new-user` and `--new-user-password` options specify the credentials for a new administrator user. The Zookeeper password is used to protect sensitive data in the Fabric registry service (all of the nodes under `/fabric`). The `--manual-ip` option specifies the Fabric Server's static IP address `StaticIPAddress` (see [the section called "Static IP address required for Fabric Server"](#)).

For more details on `fabric:create` see [section "Description" in "Console Reference" section "Description" in "Console Reference"](#).

FABRIC CREATION PROCESS

Several things happen when a fabric is created from a standalone container:

1. The container installs the requisite OSGi bundles to become a Fabric Server.
2. The Fabric Server starts a registry service, which listens on IP port 2181 (which makes fabric configuration data available to all of the containers in the fabric).
3. The Fabric Server installs a new JAAS realm (based on the ZooKeeper login module), which overrides the default JAAS realm and stores its user data in the ZooKeeper registry.
4. The new Fabric Ensemble consists of a *single* Fabric Server (the current container).
5. A default set of profiles is imported from `InstallDir/fabric/import` (can optionally be overridden).
6. After the standalone container is converted into a Fabric Server, the previously installed OSGi bundles and Karaf features are completely cleared away and replaced by the default Fabric Server configuration. For example, some of the shell command sets that were available in the standalone container are no longer available in the Fabric Server.

EXPANDING A FABRIC

You can expand a fabric by creating new managed containers. Fabric supports the *container provider* plug-in mechanism, which makes it possible to define how to create new containers in different contexts. Currently, Fabric makes container providers available for the following kinds of container:

- *Child container*, created on the local machine as a child process in its own JVM.
Instructions on creating a child container are found in [Child Containers](#).
- *SSH container*, created on any remote machine for which you have `ssh` access.
Instructions on creating a SSH container are found in [SSH Containers](#).
- *Cloud container*, created on compute instance in the cloud.
Instructions on creating a cloud container are found in [Cloud Containers](#).

Fabric provides container creation commands that make it easy to create new containers. Using these commands, Fabric can automatically install JBoss Fuse on a remote host (uploading whatever dependencies are needed), start up the remote container process, and join the container to the existing fabric, so that it becomes a fully-fledged managed container in the fabric.

CHAPTER 4. FABRIC CONTAINERS

4.1. CHILD CONTAINERS

Abstract

Child containers are the easiest kind of container to create. They are created on the same host as an existing container and are piggybacked on the same JBoss Fuse installation.

Overview

If you want to run multiple JBoss Fuse containers on a single physical host, typically the best approach is to create child containers. A child container is a relatively lightweight way to create a new container, because it re-uses most of the files in a JBoss Fuse installation. It is also convenient for administration, because the children are defined to have a parent container, so that the containers form an orderly hierarchy.

One container or many?

In principle, a single OSGi container can host multiple applications (even applications with different dependencies). So, why might you need to define extra child containers on the same host? One reason for using child containers is simply to provide a degree of isolation between applications or between components of an application. A child container runs in its own JVM process, so it is well insulated from other containers running on the same host. Using child containers also gives your application a coarse-grained structure that can be useful for managing the system (for example, each child container can be independently stopped and started).

Creating a child container

To create a new child container, invoke the `fabric:container-create-child` command, specifying the parent container name and the name of the new child container. For example, to create the new child container, `onlychild`, with `root` as its parent, enter the following command:

```
fabric:container-create-child root onlychild
```

If you want to create multiple child containers, an easy way to do this is to add an extra parameter, which specifies the number of new children. For example, to create three new child containers, enter a command like the following:

```
fabric:container-create-child root child 3
```

The preceding command would create the following new child containers:

```
child1  
child2  
child3
```

Stopping and starting a child container


Because each child container runs as a separate process, its lifecycle is independent of the parent container. That is, when you shut down a parent container, it does *not* automatically shut down the

children. To shut down a child container, you must explicitly invoke the `fabric:container-stop` command. For example, to shut down the `child1` container:

```
fabric:container-stop child1
```

To restart a stopped child container, invoke the `fabric:container-start` command, as follows:

```
fabric:container-start child1
```



NOTE

You can also stop a child container using the standard UNIX process management utilities, `ps` and `kill`.

Deleting a child container

To delete a child container (that is, permanently removing all trace of the container from the fabric, including Fabric registry entries, and data stored in the local filesystem), invoke the `fabric:container-delete` command, as follows:

```
fabric:container-delete child1
```

4.2. SSH CONTAINERS

Abstract

Fabric allows you to install containers in a local network using SSH. Fabric installs the container from scratch and configures the container to join the Fabric cluster automatically.

Overview

An SSH container is just a Fabric container that is running on a remote host on your local network, where that host is accessible through the SSH protocol. This section describes some basic administration tasks for these SSH containers.

Prerequisites

The requirements for creating an SSH container on a remote host are:

- Linux or UNIX operating system,
- SSHD running on the target host and:
 - A valid account credentials, or
 - Configured public key authentication
- Java 1.6 installed.
- Curl installed.

- GNU tar installed.
- Telnet installed.

Creating an SSH container

Fabric provides the `fabric:container-create-ssh` console command, for creating SSH containers.

Given the host, `myhost` (accessible from the local network) with the SSH user account, `myuser`, and the password, `mypassword`, you could create an SSH container on `myhost`, using the following console command:

```
fabric:container-create-ssh --host myhost --user myuser --password  
mypassword myremotecontainername
```

If the `myuser` user on `myhost` has configured public key authentication for SSH, you can skip the `password` option:

```
fabric:container-create-ssh --host myhost --user myuser  
myremotecontainername
```

Where the preceding command uses the key located in `~/.ssh/id_rsa` for authentication. If you need to use a different key, you can specify the key location explicitly with the `--private-key` option:

```
fabric:container-create-ssh --host myhost --user myuser --private-key  
~/.ssh/fabric_pk myremotecontainername
```

The last command also supports the `--pass-phrase` option, in case your key requires a pass phrase.

Creating a Fabric server using SSH

Sometimes you do not have an existing fabric and you want to create one on a remote host. The starting point for any fabric is a Fabric server instance, which can act as a seed for the rest of the fabric. So, to enable you to create a new fabric on a remote host, the `fabric:container-create-ssh` supports the `--ensemble-server` option, which can be invoked to create a container which is a Fabric server. For example, the following `container-create-ssh` command creates a new fabric consisting of one Fabric server on the `myhost` host:

```
fabric:container-create-ssh --host myhost --user myuser --ensemble-server  
myremotecontainername  
fabric:join myhost:2181
```

The `fabric:join` command joins the current container to the new fabric. This has the advantage that it is much easier to administer the new fabric using a container that is joined to the fabric, because the local container then gains access to the connection data stored in the Fabric registry.



NOTE

The argument to `fabric:join` is the ZooKeeper server port, `Host:Port`. The port number in this example, `2181`, is the standard ZooKeeper port number in Fabric.

**NOTE**

After you enter the `fabric:join` command, you will be prompted to enter the ZooKeeper password for the fabric.

Managing remote SSH containers

Using JBoss Fuse console commands, you can stop, restart or delete (that is, uninstall) a remote container, as follows:

To stop an SSH container:

```
fabric:container-stop myremotecontainername
```

To restart an SSH container:

```
fabric:container-start myremotecontainername
```

To uninstall an SSH container:

```
fabric:container-delete myremotecontainername
```

Note that these commands are available only for containers created directly using the current fabric. They are *not* available for containers that were joined to the cluster manually.

References

For more details about the SSH container console commands, see the *JBoss Fuse Console Reference*

4.3. FABRIC CONTAINERS ON WINDOWS**Abstract**

Fabric supports the deployment of containers on Windows platforms. In this case, however, it is necessary to install the container manually on the target host.

Overview

Because Windows does not support the Secure Shell (SSH) protocol, it is not possible to install the container software remotely on to a Windows machine. The installation must be performed manually. But the remote deployment of applications (by assigning profiles to the container) is fully supported.

Creating a Fabric container on Windows

Perform the following steps to create a Fabric container on Windows (assuming the container is to join an existing fabric):

1. Following the instructions in the *JBoss Fuse Installation Guide*, manually install the JBoss Fuse product on the Windows target host.
2. Open a new command prompt and enter the following commands to start the container on the target host:

```
cd InstallDir\bin
fuse.bat
```

3. If the Fabric servers from the Fabric ensemble are not already running, start them now.
4. Join the container to the existing fabric, by entering the following console command:

```
JBossFuse:karaf@root> fabric:join --zookeeper-password ZooPass
RegistryHost
```

Where **ZooPass** is the ZooKeeper password for the Fabric ensemble (as specified when you originally created the fabric with `fabric:create`); and **RegistryHost** is the hostname or IP address of one of the hosts where a Fabric server is running.



NOTE

By default, Fabric uses the default IP port number, **2181**, to connect to the Fabric server on the **RegistryHost** host. If, for some reason, the ZooKeeper service is listening on a different IP port, you can specify the IP port number explicitly using the syntax, **RegistryHost:RegistryIPPort**.

After joining the fabric, the container becomes a managed Fabric container and initially has the `default` profile deployed on it.

Creating a Fabric server on Windows

If you don't have an existing fabric, you can create a new fabric on the Windows host. The starting point for any fabric is a Fabric server instance, which can act as a seed for the rest of the fabric. Perform the following steps to create a Fabric server on Windows:

1. Following the instructions in the *JBoss Fuse Installation Guide*, manually install the JBoss Fuse product on the Windows target host.
2. To start the container on the target host, open a new command prompt and enter the following commands:

```
cd InstallDir\bin
fuse.bat
```

3. To create a new fabric (thereby turning the current host into a Fabric server), enter the following console command:

```
fabric:create --new-user AdminUser --new-user-password AdminPass
--zookeeper-password ZooPass
--resolver manualip --manual-ip StaticIPAddress --wait-for-
provisioning
```

The current container, named `root` by default, becomes a Fabric Server with a registry service installed. Initially, this is the only container in the fabric. The `--new-user` and `--new-user-password` options specify the credentials for a new administrator user. The Zookeeper

password is used to protect sensitive data in the Fabric registry service (all of the nodes under `/fabric`). The `--manual-ip` option specifies the Fabric Server's static IP address `StaticIPAddress` (see [the section called “Static IP address required for Fabric Server”](#)).

Managing remote containers on Windows

Because a Fabric container on Windows is added to the fabric by joining (that is, using `fabric:join`), there are certain restrictions on which commands you can use to manage it. In particular, the following commands are *not* supported:

```
fabric:container-stop
fabric:container-start
fabric:container-delete
```

To stop and start a Fabric container running on Windows, you must log on to the Windows host and use the regular Windows system commands to manage the container process (in particular, you could potentially install the container as a Windows service and use the Windows service commands to manage the container lifecycle).

4.4. CLOUD CONTAINERS

Abstract

Fabric has the capability to create and manage containers running in the cloud. With just a few commands, you can create a complete Fabric, consisting of multiple containers, running in a public or private cloud.

4.4.1. Preparing to use Fabric in the Cloud

Overview

Fabric leverages [JClouds](#) to enable Fabric to create new containers in public or private clouds. The Fabric cloud container provider enables you to create new compute instances in the cloud provider of your choice, perform firewall configuration, install prerequisites, install the JBoss Fuse container, and automatically register the new container.

Prerequisites

The prerequisites for creating a cloud container are as follows:

- A valid account with one of the cloud providers implemented by JClouds. The list of cloud providers can be found at [JClouds supported providers](#).



NOTE

In the context of JClouds, the term *supported provider* does not imply commercial support for the listed cloud providers. It just indicates that there is an available implementation.

Hybrid clusters

A *hybrid cluster* is a cluster composed of containers running both on the premises and on a public cloud provider. This special type of cluster has the additional requirement that *all containers must be able to connect to the Fabric registry*.

In order to satisfy this requirement, you need to make sure that one of the following conditions are met:

- Fabric registry is running inside the public cloud.

In this case, local containers will have no problem accessing the registry, as long as they are able to connect to the Internet.

- Cloud and local containers are part of a Virtual Private Network (VPN).

If the Fabric registry is running on the premises, the cloud containers will not be able to access the registry, unless you set up a VPN (or make the registry accessible from the Internet, which is *not* recommended).

- Fabric registry is accessible from the Internet (*not recommended*).

The easiest approach is to host the registry in the cloud and then configure the cloud's firewall, so that it only allows access from the containers on your premises. By default, Fabric will configure the firewall for you.

Preparation

Before you can start working with cloud containers, you must convert your local container into a Fabric container, by invoking the `fabric:create` command. You *cannot* access the requisite cloud console commands until you create a Fabric locally.

To create the Fabric container, enter the following console command:

```
JBossFuse:karaf@root> fabric:create --new-user AdminUser --new-user-
password AdminPass
--zookeeper-password ZooPass --wait-for-provisioning
```

The `--new-user` and `--new-user-password` options specify the credentials for a new administrator user. The `ZooPass` password specifies the password that is used to protect the Zookeeper registry.



NOTE

If you use a VPN (virtual private network) on your local machine, it is advisable to log off VPN *before* you create the fabric and to stay logged off while you are using the local container. A local Fabric Server is permanently associated with a fixed IP address or hostname. If VPN is enabled when you create the fabric, the underlying Java runtime is liable to detect and use the VPN hostname instead of your permanent local hostname. This can also be an issue with multi-homed machines. To be absolutely sure about the hostname, you could specify the IP address explicitly—see [Chapter 3, Creating a New Fabric](#).

The next step is to install the console commands that will enable you to administer the cloud. You can do this by adding one of the cloud profiles to your local container. The following cloud profiles are available:

```
JBossFuse:karaf@root> profile-list
[id]                [# containers] [parents]
...
cloud-aws.ec2      0                cloud-base
...
cloud-openstack    0                cloud-base
cloud-servers.uk   0                cloud-base
cloud-servers.us   0                cloud-base
...
```

For example, to install the requisite JClouds commands for interacting with the Amazon EC2 cloud, deploy the `cloud-aws.ec2` profile, as follows:

```
fabric:container-add-profile root cloud-aws.ec2
```

Where we have assumed that `root` is the name of your local container.

Feature naming convention

The most important ingredient of the `cloud-aws.ec2` profile is the `jclouds-aws-ec2` feature, which provides the necessary bundles for interacting with Amazon EC2:

```
JBossFuse:karaf@root> profile-display cloud-aws.ec2
Profile id: cloud-aws.ec2
Version   : 1.0
...
Container settings
-----
Features :
  jclouds-aws-ec2
...
```

Some commonly used cloud providers can be accessed using the following Karaf features:

`jclouds-aws-ec2`

Feature for the Amazon EC2 cloud provider.

`jclouds-cloudservers-us`

Feature for the Rackspace cloud provider.

In general, the naming convention for cloud provider features is: `jclouds-ProviderID`, where *ProviderID* is one of the provider IDs listed in the [JClouds supported providers](#) page. Or you can list the available JClouds features using the `features:list` command:

```
features:list | grep jclouds
```

If you want to add another JClouds feature to your container, add it to a Fabric profile and then deploy the profile to your container (or add the feature to a profile that is already deployed). For example:

```
fabric:profile-edit --features jclouds-ProviderID MyProfile
fabric:container-add-profile root MyProfile
```

Registering a cloud provider

After installing the required cloud features, you need to register the cloud provider with Fabric, using the `fabric:cloud-service-add` console command (the registration process will store the provider credentials in the Fabric registry, so that they are available from any Fabric container).

You need to obtain a valid *identity* and *credential* from your cloud provider, which are not necessarily the same thing as the username and password you obtained upon registration with the provider. Usually, they refer to the credentials you get for using the cloud service from an external API. For example, on Amazon EC2 the requisite credentials can be found on the [security credentials page](#).

For example, to register the Amazon EC2 provider:

```
fabric:cloud-service-add --name aws-ec2 --provider aws-ec2
--identity AccessKeyID --credential SecretAccessKey
```



NOTE

The identifier supplied to the `--name` option is an alias that you use to refer to this registered cloud provider instance. It is possible to register the same cloud provider more than once, with different user accounts. The cloud provider alias thus enables you to distinguish between multiple accounts with the same cloud provider.

4.4.2. Administering Cloud Containers

Creating a new fabric in the cloud

To create a fabric in the cloud, invoke the `fabric:container-create-cloud` with the `--ensemble-server` option, which creates a new Fabric server. For example, to create a Fabric server on Amazon EC2:

```
fabric:container-create-cloud --ensemble-server --name aws-ec2
--new-user AdminUser --new-user-password AdminPass --zookeeper-password
ZooPass mycontainer
```

Basic security

When creating a new fabric in the cloud, it is necessary to supply some basic security information to the `fabric:container-create-cloud` command, to ensure that the new fabric is adequately protected. You need to specify the following security data:

- *JAAS credentials*—the `--new-user` and `--new-user-password` options define JAAS credentials for a new user with administrative privileges on the fabric. These credentials can subsequently be used to log on to the JMX port or the SSH port of the newly created Fabric server.
- *ZooKeeper password*—is used to protect the data stored in the ZooKeeper registry in the Fabric server. The only time you will be prompted to enter the ZooKeeper password is when you try to join a container to the fabric using the `fabric:join` command.

Joining a standalone container to the fabric

If you have been using a standalone container (not part of a fabric) to create the fabric in the cloud, it is

a good idea to join this container to the newly created fabric, so that you can easily administer the fabric from your local container. To join your local container to the fabric, enter a command like the following:

```
fabric:join -n --zookeeper-password ZooPass PublicIPAddress
```

Where *PublicIPAddress* is the public host name or the public IP address of the compute instance that hosts the Fabric server (you can get this address either from the JBoss Fuse console output or from the Amazon EC2 console).

Alternatively, instead of joining your local container to the fabric, you could use the JBoss Fuse client utility to log into the remote Fabric server directly (using the JAAS credentials).

Creating a cloud container

After creating the initial Fabric server (which constitutes the Fabric ensemble), you can use the `fabric:container-create-cloud` command to create new Fabric containers in the cloud. For example to create a container on Amazon EC2:

```
fabric:container-create-cloud --name aws-ec2 mycontainer
```

Specifying an image is optional. By default, Fabric tries to find an Ubuntu image for you. You can provide options for the operating system and the O/S version. For example, to choose Centos instead of Ubuntu, you could invoke the `fabric:container-create-cloud` command with the `--os-family` option as follows:

```
fabric:container-create-cloud --name aws-ec2 --os-family centos
mycontainer
```

Or to be even more specific, you can specify the O/S version as well, using the `--os-version` option:

```
fabric:container-create-cloud --name aws-ec2 --os-family centos --os-
version 5 mycontainer
```

If you need to specify the exact image, use the `--image` option.

```
fabric:container-create-cloud --name aws-ec2 --image myimageid mycontainer
```

After creating the new cloud container, the command displays the creation status and some useful information:

```
Looking up for compute service.
Creating 1 nodes in the cloud. Using operating system: ubuntu. It may take
a while ...
Node fabric-f674a68f has been created.
Configuring firewall.
Installing fabric agent on container cloud. It may take a while...
Overriding resolver to publichostname.
                                [id] [container]                                [public addresses]
[status]
  us-east-1/i-f674a68f cloud                                [23.20.114.82]
success
```

Images

Regardless of the way that you specify the image (directly or indirectly), the image needs to have some of the following characteristics:

- Linux O/S
- RedHat or Debian packaging style
- Either no Java installed or Java 1.6+ installed. If there is no Java installed on the image, Fabric will install Java for you. If the wrong Java version is installed, however, the container installation will fail.

If you prefer, you can create your own custom image and use that instead. But this typically requires some additional configuration when you register the cloud provider. For example, on Amazon EC2 you would need to specify the owner ID of the private image when registering the provider:

```
fabric:cloud-service-add --name aws-ec2 --provider aws-ec2  
--identity AccessKeyID --credential SecretAccessKey --owner myownerid
```

Locations and hardware

Most cloud providers will give you the option to create containers on different locations or using different hardware profiles. You may wonder which are the proper values to use for your provider. Even though Fabric provides completion for **all** configuration options, you still may want to get a list of them.

To list all of the available locations:

```
jclouds:location-list
```

To list all the available hardware profiles:

```
jclouds:hardware-list
```

To exploit this information for creating a cloud container, you can specify them as options to the **fabric:container-create-cloud** command. For example:

```
fabric:container-create-cloud --name aws-ec2 --location eu-west-1 --  
hardware m2.4xlarge mycontainer
```


CHAPTER 5. FABRIC PROFILES

Abstract

A profile is the basic unit of deployment in a fabric. This chapter describes how to create, edit, and deploy profiles into containers. You can also create different versions of a profile, which makes it possible to support rolling upgrades across the containers in your fabric.

5.1. INTRODUCTION TO PROFILES

Overview

A profile is a description of how to provision a logical group of containers. Each profile can have none, one, or more parents, which allows you to have profile hierarchies. A container can be assigned one or more profiles. Profiles are also versioned, which enables you to maintain different versions of each profile, and then upgrade or roll back containers, by changing the version of the profiles they use.

What is in a profile?

A profile can contain one or more of the following resources:

- OSGi bundle URLs
- Web ARchive (WAR) URLs
- Fuse Application Bundle (FAB) URLs
- OSGi Configuration Admin PIDs
- Apache Karaf feature repository URLs
- Apache Karaf features
- Maven artifact repository URLs
- Blueprint XML files or Spring XML files (for example, for defining broker configurations or Camel routes)
- Any kind of resource that might be needed by an application (for example, Java properties file, JSON file, XML file, YML file)
- System properties that affect the Apache Karaf container (analogous to editing `etc/config.properties`)
- System properties that affect installed bundles (analogous to editing `etc/system.properties`)

Profile hierarchies

Frequently, multiple profiles share a lot of configuration details: such as common frameworks, libraries, and so on. Defining these details separately for each profile would create a considerable maintenance headache. To avoid duplication across profiles, therefore, Fabric uses a hierarchical model for profiles.

You can define a generic profile (base profile) containing the common configuration details, and then define child profiles that inherit these generic configuration details.

Some basic profiles

Fabric provides a rich set of predefined profiles, which can be used as the basic building blocks for defining your own profiles. Some of the more interesting predefined profiles are:

[default]

The `default` profile defines all of the basic requirements for a Fabric container. For example it specifies the `fabric-agent` feature, the Fabric registry URL, and the list of Maven repositories from which artifacts can be downloaded.

[karaf]

Inherits from the `default` profile and defines the Karaf feature repositories, which makes the Apache Karaf features accessible.

[feature-camel]

Inherits from `karaf`, defines the Camel feature repositories, and installs some core Camel features: such as `camel-core` and `camel-blueprint`. If you are deploying a Camel application, it is recommended that you inherit from this profile.

[feature-cxf]

Inherits from `karaf`, defines the CXF feature repositories, and installs some core CXF features. If you are deploying a CXF application, it is recommended that you inherit from this profile.

[mq-base]

Inherits from the `karaf` profile and installs the `mq-fabric` feature

[mq-default]

Inherits from the `mq-base` profile and provides the configuration for an A-MQ broker. Use this profile, if you want to deploy a minimal installation of an ActiveMQ broker.

[jboss-fuse-minimal]

Includes all of the features and bundles required for the JBoss Fuse minimal container.

[jboss-fuse-medium]

Includes all of the features and bundles required for the JBoss Fuse medium container.

[jboss-fuse-full]

Includes all of the features and bundles required for the JBoss Fuse full container.

5.2. WORKING WITH PROFILES

Changing the profiles in a container

To change the profiles assigned to a Fabric container, invoke the `fabric:container-change-profile` command as follows:

```
fabric:container-change-profile mycontainer myprofile
```

Where the preceding command deploys the **myprofile** profile to the **mycontainer** container. All profiles previously assigned to the container are removed. You can also deploy multiple profiles to the container, with the following command:

```
fabric:container-change-profile mycontainer myprofile myotherprofile
```

Adding a profile to a container

The **fabric:container-add-profile** command gives you a simple way to add profiles to a container, without having to list all of the profiles that were already assigned. For example, to add the **example-camel** profile to the **mycontainer** container:

```
fabric:container-add-profile mycontainer example-camel
```

Listing available profiles

To see the list of available profiles, invoke the **fabric:profile-list** console command:

```
fabric:profile-list
```

The command displays all available profiles, showing their parents and the number of containers each profile is deployed into.

Inspecting profiles

To see exactly what a profile defines, enter the **fabric:profile-display** command. For example, to display what is defined in the **feature-camel** profile, enter the following command:

```
fabric:profile-display feature-camel
```

Which outputs something like the following to the console window:

```
Profile id: feature-camel
Version   : 1.0
Attributes:
  parents: karaf

Containers:

Container settings
-----
Repositories :
  mvn:org.apache.camel.karaf/apache-camel/${version:camel}/xml/features

Features :
  camel-core
  camel-blueprint
  fabric-camel

Configuration details
```

```

-----
Other resources
-----
Resource: org.fusesource.insight.metrics.json

```

The preceding output does not take into account the definitions inherited from any parent profiles, however. To see the effective definitions for the `feature-camel` profile, taking into account all of its ancestors, you must specify the `--overlay` switch, as follows:

```
fabric:profile-display --overlay feature-camel
```

Resource files stored in the profile are listed under the heading `Other resources`. If you want to display the contents of these resource files as well, add the `--display-resources` switch (or `-r` for short) to the `profile-display` command, as follows:

```
fabric:profile-display -r feature-camel
```

Creating a new profile

To create a new profile for an application, invoke the `fabric:profile-create` command, as follows:

```
fabric:profile-create myprofile
```

To specify one or more parents for the profile when it is being created, add the `--parents` option to the command:

```
fabric:profile-create --parents feature-camel myprofile
```

After the profile is created, you can start to modify the profile, providing details of what should be deployed in the profile.

Adding or removing features

To edit one of the existing profiles, you can use the `fabric:profile-edit` command. For example, to add the `camel-jclouds` feature to the `feature-camel` profile.

```
fabric:profile-edit --features camel-jclouds feature-camel
```

Now invoke the `fabric:profile-display` command to see what the camel profile looks like now. You should see that the `camel-jclouds` feature appears in the list of features for the `feature-camel` profile.

```

Features :
  camel-jclouds
  camel-blueprint/2.9.0.fuse-7-061
  camel-core/2.9.0.fuse-7-061
  fabric-camel/99-master-SNAPSHOT

```

If you want to remove a feature from the profile, use the `--delete` option. For example, if you need to remove the `camel-jclouds` feature, you could use the following command:

```
fabric:profile-edit --delete --features camel-jclouds feature-camel
```

Editing PID properties

An OSGi Config Admin Persistent ID (PID) consists essentially of a list of key-value pairs. You can edit PID properties using either of the following approaches:

- *Edit the PID using the built-in text editor*—the Karaf console has a built-in text editor which you can use to edit profile resources such as PID properties. To start editing a PID using the text editor, enter the following console command:

```
fabric:profile-edit --pid PID ProfileName
```

For more details about the built-in text editor, see [Appendix A, Editing Profiles with the Built-In Text Editor](#).

- *Edit the PID inline, using console commands*—alternatively, you can edit PIDs directly from the console, using the appropriate form of the `fabric:profile-edit` command. This approach is particularly useful for scripting. For example, to set a specific key-value pair, *Key=Value*, in a PID, enter the following console command:

```
fabric:profile-edit --pid PID/Key=Value ProfileName
```

Editing a PID inline

To edit a PID inline, use the following variants of the `fabric:profile-edit` command:

- Assign a value to a PID property, as follows:

```
fabric:profile-edit --pid PID/Key=Value ProfileName
```

- Append a value to a delimited list (that is, where the property value is a *comma-separated list*), as follows:

```
fabric:profile-edit --append --pid PID/Key=ListItem ProfileName
```

- Remove a value from a delimited list, as follows:

```
fabric:profile-edit --remove --pid PID/Key=ListItem ProfileName
```

- Delete a specific property key, as follows:

```
fabric:profile-edit --delete --pid PID/Key ProfileName
```

- Delete a complete PID, as follows:

```
fabric:profile-edit --delete --pid PID ProfileName
```

Example of editing a PID inline

In the following example, we modify the `io.fabric8.agent` PID, changing the Maven repository list setting. The `default` profile contains a section like this:

```
Agent Properties :
    org.ops4j.pax.url.mvn.repositories =
http://repo1.maven.org/maven2,

http://repo.fusesource.com/nexus/content/repositories/releases,
    http://repo.fusesource.com/nexus/content/groups/ea,
    http://repository.springsource.com/maven/bundles/release,
    http://repository.springsource.com/maven/bundles/external,
    http://scala-tools.org/repo-releases
```

The agent properties section is represented by the `io.fabric8.agent` PID. So, by modifying the `io.fabric8.agent` PID, we effectively change the agent properties. You can modify the list of Maven repositories in the agent properties PID as follows:

```
fabric:profile-edit --pid
io.fabric8.agent/org.ops4j.pax.url.mvn.repositories=http://repositorymanag
er.mylocalnetwork.net default
```

Now when you invoke `fabric:profile-display` on the `default` profile, you should see agent properties similar to the following:

```
Agent Properties :
    org.ops4j.pax.url.mvn.repositories =
http://repositorymanager.mylocalnetwork.net
```

Setting encrypted PID property values

In some cases, you might prefer to store PID property values in encrypted format instead of plain text. For example, passwords and other sensitive data should usually be stored in encrypted form. To store a property value in encrypted form, perform the following steps:

1. Use the `fabric:encrypt-message` command to encrypt the property value, as follows:

```
fabric:encrypt-message PropValue
```

This command returns the encrypted property value, *EncryptedValue*.



NOTE

The default encryption algorithm used by Fabric is `PBEWithMD5AndDES`.

2. You can now set the property to the encrypted value, *EncryptedValue*, using the following syntax:

```
my.sensitive.property = ${crypt:EncryptedValue}
```

For example, using the `fabric:profile-edit` command, you can set an encrypted value as follows:

■

```
fabric:profile-edit --pid
com.example.my.pid/my.sensitive.property=${crypt:EncryptedValue}
Profile
```



WARNING

These encrypted values are protected by the master password, which is accessible to anyone who can log on to a Fabric container. To keep these encrypted values safe, you must restrict access to the containers in the fabric.

Alternative method for encrypting PID property values

The underlying encryption mechanism for PID properties is based on the [Jasypt](#) encryption toolkit. Consequently, it is also possible to encrypt PID properties directly, using the Jasypt toolkit, as follows:

1. [Download and install Jasypt](#), to gain access to the Jasypt `encrypt` and `decrypt` command-line tools.
2. Use the Jasypt `encrypt` command-line tool to encrypt the property value, as follows:

```
./encrypt.sh input="Property value to be encrypted" password=ZooPass
verbose=false
```

This command returns the encrypted property value, *EncryptedValue*.



NOTE

The default encryption algorithm used by Fabric is **PBEwithMD5andDES**. You must ensure that the `encrypt.sh` utility is using the same algorithm as Fabric.

Customizing the PID property encryption mechanism

You can customize the PID property encryption mechanism, as follows:

- *Customize the master password for encryption*—using the following console command:

```
fabric:crypt-password-set MasterPassword
```

You can retrieve the current master password by entering the `fabric:crypt-password-get` command. The default value is the ensemble password (as returned by `fabric:ensemble-password`).

- *Customize the encryption algorithm*—using the following console command:

```
fabric:crypt-algorithm-set Algorithm
```

Where the encryption algorithm must be one of the algorithms supported by the underlying [Jasypt](#) encryption toolkit. You can retrieve the current encryption algorithm by entering the `fabric:crypto-algorithm-get` command. The default is `PBEwithMD5AndDES`.

Profile editor

If you want to make extensive edits to a profile, it is not very convenient to make changes one setting at a time. There is a more convenient approach for making extensive profile edits, and that is to use the console's built-in profile editor, which is a simple screen-based text editor.

For example, to open the agent properties resource for editing, simply invoke the `fabric:profile-edit` command without any options, as follows:

```
fabric:profile-edit Profile [Version]
```

A simple text editor opens, enabling to edit the configuration settings in the agent properties.

For full details of how to edit profiles using the built-in text editor, see [Appendix A, Editing Profiles with the Built-In Text Editor](#).

Editing resources with the profile editor

A practical way to edit a general profile resource (such as an XML configuration resource) is to use the built-in text editor. For example, to start editing the `broker.xml` file in the `mq-amq` profile, enter the following console command:

```
fabric:profile-edit --resource broker.xml mq-amq
```

5.3. PROFILE VERSIONS

Overview

Every profile has at least one version. When assigning a profile to a container, you actually assign both the profile and the version. The `fabric-agent`, will choose the defined version and retrieve all the information provided by the specific version of the profile.

Any change to a profile takes immediate effect. This means that any container using a profile that was just modified will pick up the change immediately. It is recommended that you create a new version of a profile whenever you need to make changes. You can then upgrade containers to use the new version. This enables you to perform atomic updates, test updates on specific containers, and possibly roll back to the previous version, if you encounter any problems.

Creating a new version

You can create a new version using the `fabric:version-create` command (analogous to creating a new branch in the underlying Git repository). The default version is 1.0. To create version 1.1, enter the following command:

```
fabric:version-create 1.1
```

After the 1.1 version is created, a new instance of every profile is created for the new version (copied from the previous latest version, which was 1.0). Now you can display or modify the 1.1 version of each

profile. For example, enter the following command to display the details of the `feature-camel` profile:

```
fabric:profile-display --version 1.1 feature-camel
```

Initially, the output is identical to the 1.0 version of the profile, because we have not yet modified the new version of the profile. But how do you modify a specific version of a profile? All you need to do is to invoke the `fabric:profile-edit` command, specifying the version right after the profile argument. For example, to add the `camel-jclouds` feature to version 1.1 of the `feature-camel` profile, enter the following command:

```
fabric:profile-edit --features camel-jclouds feature-camel 1.1
```



IMPORTANT

The changes made to version 1.1 of the profile do *not* (yet) affect any of your existing containers. The changes do not take effect until you upgrade your containers to use the 1.1 version.

Rolling upgrades and rollbacks

Fabric provides commands for *upgrading* (incrementing the effective version) and *rolling back* (decrementing the effective version) the profiles assigned to a container. For example, to upgrade the `mycontainer` container to the 1.1 version, invoke the `fabric:container-upgrade` command as follows:

```
fabric:container-upgrade 1.1 mycontainer
```

The preceding command makes `mycontainer` to use version 1.1 of all the profiles currently assigned to it.

If for any reason you want to roll back to the previous version, you can invoke the `fabric:container-rollback` command, as follows:

```
fabric:container-rollback 1.0 mycontainer
```

It is strongly recommended that you test any profile changes on a *single* container, before applying the changes to the whole cluster. Applying an upgrade to all containers can be achieved by specifying the `--all` option, as follows:

```
fabric:container-upgrade --all 1.1 mycontainer
```

CHAPTER 6. FABRIC8 MAVEN PLUG-IN

Abstract

This Maven plug-in makes it easy to create or update a fabric profile from your Maven project.



WARNING

The Fabric8 Maven plug-in is a technical preview only, and is *not* supported in JBoss A-MQ 6.1.

6.1. PREPARING TO USE THE PLUG-IN

Edit your Maven settings

First you will need to edit your `~/.m2/settings.xml` file to add the fabric server's user and password so that the maven plugin can log in to the fabric. For example, you could add the following `server` element to your `settings.xml` file:

```
<settings>
  <servers>
    <server>
      <id>fabric8.upload.repo</id>
      <username>Username</username>
      <password>Password</password>
    </server>
    ...
  </servers>
</settings>
```

Where *Username* and *Password* are the credentials of a Fabric user with administrative privileges (for example, the credentials you would use to log on to the Management Console).

Customising the repository ID

The default Fabric Maven repository ID is `fabric8.upload.repo`. You can specify additional `server` elements in your `settings.xml` file for each of the fabrics you need to work with. To select the relevant credentials, you can set the `serverId` property in the Fabric8 Maven plug-in `configuration` section (see [Section 6.4, “Configuration Properties”](#)) or set the `fabric8.serverId` Maven property.

6.2. USING THE PLUG-IN TO DEPLOY A MAVEN PROJECT

Prerequisites

You must ensure the following prerequisites are satisfied before attempting to run the Fabric8 Maven

plug-in:

1. Your Maven `~/.m2/settings.xml` file is configured as described in [Section 6.1, “Preparing to Use the Plug-In”](#).
2. A JBoss Fuse container instance is running on your local machine (alternatively, if the container instance is running on a remote host, you must configure the plug-in's `jolokiaUrl` property appropriately).

Running the plug-in on any Maven project

To use the Fabric8 plug-in to deploy any maven project into a fabric profile, enter the following Maven command:

```
mvn io.fabric8:fabric8-maven-plugin:1.0.0.redhat-355:deploy
```

Adding the plug-in to a Maven POM

If you add the Fabric8 plug-in to your `pom.xml` file as follows:

```
<plugins>
  <plugin>
    <groupId>io.fabric8</groupId>
    <artifactId>fabric8-maven-plugin</artifactId>
  </plugin>
</plugins>
```

You can use the following more concise Maven goal:

```
mvn fabric8:deploy
```

What does the plug-in do?

When you deploy your project to a Fabric profile with this plug-in, the plug-in does the following:

- Uploads any artifacts into the fabric's Maven repository,
- Lazily creates the Fabric profile or version you specify,
- Adds/updates the Maven project artifact into the profile configuration,
- Adds any additional parent profile, bundles or features to the profile.

Example

You can try out the plug-in with one of the JBoss Fuse `quickstart` examples, as follows:

```
cd $INSTALLDIR/quickstarts/rest
mvn io.fabric8:fabric8-maven-plugin:1.0.0.redhat-355:deploy
```

You should see a new profile created at the [my-rest/rest profile page](#), which should have a bundle and some features (click on the **Bundle** tab and the **Feature** tab).

6.3. CONFIGURING THE PLUG-IN

Specifying profile information

You can explicitly configure the name of the profile to create, by adding a `configuration` element to the plug-in configuration in your `pom.xml` file, as follows:

```
<plugins>
  <plugin>
    <groupId>io.fabric8</groupId>
    <artifactId>fabric8-maven-plugin</artifactId>
    <configuration>
      <profile>my-thing</profile>
    </configuration>
  </plugin>
</plugins>
```

Multi-module Maven projects

For multi-module Maven projects, a more flexible way to configure the plug-in is to use Maven properties. For example if you have a multi-module Maven project such as this:

```
pom.xml
foo/
  pom.xml
  a/pom.xml
  b/pom.xml
  ...
bar/
  pom.xml
  c/pom.xml
  d/pom.xml
  ...
```

You could define the plug-in once in the root `pom.xml` file, as follows:

```
<plugins>
  <plugin>
    <groupId>io.fabric8</groupId>
    <artifactId>fabric8-maven-plugin</artifactId>
  </plugin>
</plugins>
```

While in the `foo/pom.xml` file you need only define the `fabric8.profile` property, as follows:

```
<project>
  ...
  <properties>
    <fabric8.profile>my-foo</fabric8.profile>
    ...
  </properties>
  ...
</project>
```

All of the projects within the `foo` folder, such as `foo/a` and `foo/b`, will deploy to the same profile (in this case the profile, `my-foo`). You can use the same approach to put all of the projects under the `bar` folder into a different profile too.

At any point in your tree of Maven projects you can define a Maven `fabric8.profile` property to specify exactly where it gets deployed; along with any other property on the plug-in (see the Property Reference below).

Specifying features, additional bundles, repositories and parent profiles

You can specify additional configuration in the maven plug-in, as follows:

```
<plugins>
  <plugin>
    <groupId>io.fabric8</groupId>
    <artifactId>fabric8-maven-plugin</artifactId>
    <configuration>
      <profile>my-rest</profile>
      <features>fabric-cxf-registry fabric-cxf cxf war swagger</features>
      <featureRepos>mvn:org.apache.cxf.karaf/apache-
cxf/${version:cxf}/xml/features</featureRepos>
    </configuration>
  </plugin>
</plugins>
```

Note that the `features` element allows you to specify a space-separated list of features to include in the profile.

This example specifies space-separated lists for the parent profile IDs, features, repositories and bundles so that it is easy to reuse Maven properties for these values (for example, to add some extra features to a child maven project while inheriting from the parent project).

Configuring with Maven properties

You can also use Maven property values (or command line arguments) to specify the configuration values by prefixing the property name with `fabric8..` For example, to deploy a Maven project to the `cheese` profile name, enter the command:

```
mvn fabric8:deploy -Dfabric8.profile=cheese -Dfabric8.version=1.1
```

By default, the project artifacts are uploaded to the Maven repository inside the fabric. If you want to disable this behavior and just update the profile configuration (for example, if you are already pointing your fabric's Maven repository to your local Maven repository), you can set `fabric8.upload=false`—for example:

```
mvn fabric8:deploy -Dfabric8.upload=false
```

Specifying profile resources

If you create the directory, `src/main/fabric8`, in your Maven project and add any resource files or a `ReadMe.md` file to your project, they will automatically be uploaded into the profile as well. For example, if you run the following commands from your Maven project directory:

```
mkdir -p src/main/fabric8
echo "## Hello World" >> src/main/fabric8/ReadMe.md
mvn fabric8:deploy
```

The newly deployed profile will include a `ReadMe.md` wiki page.

6.4. CONFIGURATION PROPERTIES

Specifying properties

Properties can be specified either as elements inside the `configuration` element of the plug-in in your project's `pom.xml` file. For example, the `profile` property can be set as follows:

```
<plugins>
  <plugin>
    <groupId>io.fabric8</groupId>
    <artifactId>fabric8-maven-plugin</artifactId>
    <configuration>
      <profile>${fabric8.profile}</profile>
    </configuration>
  </plugin>
</plugins>
```

Or you can specify the properties on the command line or as Maven build properties (where the property names must be prefixed with `fabric8..` For example, to set the profile name, you could add the following property to your `pom.xml` file:

```
<project>
  ...
  <properties>
    <fabric8.profile>my-foo</fabric8.profile>
    ...
  </properties>
  ...
```

Or you can specify properties on the command line:

```
mvn fabric8:deploy -Dfabric8.profile=my-foo
```

Property reference

The Fabric8 Maven plug-in supports the following properties (which can be set either as elements inside the `configuration` element in the `pom.xml` file or as Maven properties, when prefixed by `fabric8.`):

Parameter	Description
<code>profile</code>	The name of the Fabric profile to deploy your project to. Defaults to the <i>groupId-artifactId</i> of your Maven project.

serverId	The server ID used to lookup in <code>~/.m2/settings/xml</code> for the <code>server</code> element to find the username and password to log in to the fabric. Defaults to <code>fabric8.upload.repo</code> .
jolokiaUrl	The Jolokia URL of the JBoss Fuse Management Console. Defaults to <code>http://localhost:8181/jolokia</code> .
version	The Fabric version in which to update the profile. Defaults to the current version of the fabric.
baseVersion	If the version does not exist, the <code>baseVersion</code> provides the initial values for the newly created version. This is like creating a branch from the <code>baseVersion</code> for a new version branch in git.
parentProfiles	Space-separated list of parent profile IDs to be added to the newly created profile. Defaults to <code>karaf</code> .
features	Space-separated list of features to add to the profile. For example, the following setting would include both the <code>camel</code> feature and the <code>cxfr</code> feature: <code><features>camel cxfr</features></code>
featureRepos	Space-separated list of feature repository URLs to add to the profile. The URL has the general form <code>mvn:groupId/artifactId/version/xml/features</code> .
bundles	Space-separated list of additional bundle URLs (of the form <code>mvn:groupId/artifactId/version</code>) to add to the newly created profile. Note you do not have to include the current Maven project artifact; this configuration is intended as a way to list dependent required bundles.
upload	Whether or not the <code>deploy</code> goal should upload the local builds to the fabric's Maven repository. You can disable this step if you have configured your fabric's Maven repository to reuse your local maven repository. Defaults to <code>true</code> .
profileConfigDir	The folder in your maven project containing resource files to be deployed into the profile, along with the artifact configuration. Defaults to <code>src/main/fabric8</code> . You should create the directory and add any configuration files or documentation you wish to add to your profile.

CHAPTER 7. ACTIVEMQ BROKERS AND CLUSTERS

Abstract

Fabric provides predefined profiles for deploying a simple standalone broker and, in addition, you can use the powerful `fabric:mq-create` command to create and deploy clusters of brokers.

7.1. CREATING A STANDALONE BROKER INSTANCE

MQ profiles

The following profiles are important for creating broker instances:

`mq-base`

An abstract profile, which defines some important properties and resources for the broker, but should never be used directly to instantiate a broker.

`mq-default`

A basic standalone broker, which inherits most of its properties from the `mq-base` profile.

To examine the properties defined in these profiles, you can invoke the `fabric:profile-display` command, as follows:

```
JBossFuse:karaf@root> fabric:profile-display mq-default
...
JBossFuse:karaf@root> fabric:profile-display mq-base
...
```

Creating a new broker instance

A Fuse MQ broker is a Karaf container instance running a message broker profile. The profile defines the broker dependencies (through features) and the configuration for the broker. The simplest approach to creating a new broker is to use the provided `mq-default` profile.

For example, to create a new `mq-default` broker instance called `broker1`, enter the following console command:

```
JBossFuse:karaf@root>fabric:container-create-child --profile mq-default
root broker1
The following containers have been created successfully:
    broker1
```

This command creates a new container called `broker1` with a broker *of the same name* running on it.

`fabric:mq-create` command

The `fabric:mq-create` command provides a short cut to creating a broker, but with more flexibility, because it also creates a new profile. To create a new broker instance called `brokerx` using `fabric:mq-create`, enter the following console command:


```
JBossFuse:karaf@root> fabric:mq-create --create-container broker --
replicas 1 brokerx
MQ profile mq-broker-default.brokerx ready
```

Just like the basic `fabric:container-create-child` command, `fabric:mq-create` creates a container called `broker1` and runs a broker instance on it. There are some differences, however:

- The new `broker1` container is implicitly created as a child of the current container,
- The new broker has its own profile, `mq-broker-default.brokerx`, which is based on the `mq-base` profile template,
- It is possible to edit the `mq-broker-default.brokerx` profile, to customize the configuration of this new broker.
- The `--replicas` option lets you specify the number of master/slave broker replicas (for more details, see [Section 7.3.2, “Master-Slave Cluster”](#)). In this example, we specify one replica (the default is two).



NOTE

The new profile gets the name `mq-broker-Group.BrokerName` by default. If you want the profile to have the same name as the broker (which was the default in JBoss A-MQ version 6.0), you can specify the profile name explicitly using the `--profile` option.

Starting a broker on an existing container

The `fabric:mq-create` command can be used to deploy brokers on existing containers. Consider the following example, which creates a new Fuse MQ broker in two steps:

```
JBossFuse:karaf@root> fabric:container-create-child root broker1
The following containers have been created successfully:
  broker1

JBossFuse:karaf@root> fabric:mq-create --assign-container broker1 brokerx
MQ profile mq-broker-default.brokerx ready
```

The preceding example firstly creates a default child container, and then creates and deploys the new `mq-broker-default.brokerx` profile to the container, by invoking `fabric:mq-create` with the `--assign-container` option. Of course, instead of deploying to a local child container (as in this example), we could assign the broker to an SSH container or a cloud container.

Broker groups

Brokers created using the `fabric:mq-create` command are always registered with a specific *broker group*. If you do not specify the group name explicitly at the time you create the broker, the broker gets registered with the `default` group by default.

If you like, you can specify the group name explicitly using the `--group` option of the `fabric:mq-create` command. For example, to create a new broker that registers with the `west-coast` group, enter the following console command:

```
JBossFuse:karaf@root> fabric:mq-create --create-container broker --
replicas 1 --group west-coast brokery
MQ profile mq-broker-default.brokery ready
```

If the `west-coast` group does not exist prior to running this command, it is automatically created by Fabric. Broker groups are important for defining clusters of brokers, providing the underlying mechanism for creating load-balancing clusters and master-slave clusters. For details, see [Section 7.3, “Topologies”](#).

7.2. CONNECTING TO A BROKER

Overview

This section describes how to connect a client to a broker. In order to connect to a JBoss MQ broker, you need to know its *group name*. Every MQ broker is associated with a group when it is created: if none is specified explicitly, it automatically gets associated with the `default` group.

Client URL

To connect to an MQ broker, the client must specify a discovery URL, in the following format:

```
discovery:(fabric:GroupName)
```

For example, to connect to a broker associated with the `default` group, the client would use the following URL:

```
discovery:(fabric:default)
```

The connection factory then looks for available brokers in the group and connects the client to one of them.

Example client profiles

You can test broker by deploying the `example-mq` profile into a container. The `example-mq` profile instantiates a pair of messaging clients: a *producer client*, that sends messages continuously to the `FABRIC.DEMO` queue on the broker; and a *consumer client*, that consumes messages from the `FABRIC.DEMO` queue.

Create a new container with the `example-mq` profile, by entering the following command:

```
JBossFuse:karaf@root> fabric:container-create-child --profile example-mq
root example
```

You can check whether the `example` container is successfully provisioned, using the following console command:

```
JBossFuse:karaf@root> watch container-list
```

After the `example` container is successfully provisioned, you can connect to it and check its log to verify the flow of messages, using the following console commands:

```
JBossFuse:karaf@root> container-connect example
JBossFuse:karaf@example> log:display
```

7.3. TOPOLOGIES


7.3.1. Load-Balancing Cluster

Overview

Fabric exploits the concept of broker groups to implement cluster functionality. To set up a load-balancing cluster, all of the brokers in the cluster should register with the same group name, but using *unique* broker names.

For example, [Figure 7.1, “Load-Balancing Cluster”](#) shows a load-balancing cluster with the group name, **loadba1**, and with three brokers registered in the group: **brokerx**, **brokery**, and **brokerz**. This type of topology is ideal for load balancing non-persistent messages across brokers and for providing high-availability.

Figure 7.1. Load-Balancing Cluster



Create brokers in a load-balancing cluster

The basic rules for creating a load-balancing cluster of brokers are as follows:

- Choose a group name for the load-balancing cluster.
- Each broker in the cluster registers with the chosen group.
- Each broker must be identified by a *unique broker name*.
- Normally, each broker is deployed in a separate container.

For example, consider the cluster shown in [Figure 7.1, “Load-Balancing Cluster”](#). The group name is **loadba1** and the cluster consists of three broker instances with broker names: **brokerx**, **brokery**, and **brokerz**.

To create this cluster, perform the following steps:

1. First of all create some containers:

```
JBossFuse:karaf@root> container-create-child root broker 3
The following containers have been created successfully:
```

```

Container: broker1.
Container: broker2.
Container: broker3.

```

2. Wait until the containers are successfully provisioned. You can conveniently monitor them using the `watch` command, as follows:

```
JBossFuse:karaf@root> watch container-list
```

3. You can then assign broker profiles to each of the containers, using the `fabric:mq-create` command, as follows:

```

JBossFuse:karaf@root> mq-create --group loadbal --assign-container
broker1 brokerx
MQ profile mq-broker-loadbal.brokerx ready

JBossFuse:karaf@root> mq-create --group loadbal --assign-container
broker2 brokery
MQ profile mq-broker-loadbal.brokery ready

JBossFuse:karaf@root> mq-create --group loadbal --assign-container
broker3 brokerz
MQ profile mq-broker-loadbal.brokerz ready

```

4. You can use the `fabric:profile-list` command to see the new profiles created for these brokers:

```

JBossFuse:karaf@root> profile-list --hidden
[id]                [# containers] [parents]
...
mq-broker-loadbal.brokerx      1          mq-base
mq-broker-loadbal.brokery      1          mq-base
mq-client-loadbal
...

```

5. You can use the `fabric:cluster-list` command to view the cluster configuration for this load balancing cluster:

```

JBossFuse:karaf@root> cluster-list
[cluster]           [masters]
[slaves]           [services]
...
fusemq/loadbal
  brokerx           broker1      -
tcp://MyLocalHost:50394
  brokery           broker2      -
tcp://MyLocalHost:50604
  brokerz           broker3      -
tcp://MyLocalHost:50395

```

Configure clients of a load-balancing cluster

To connect a client to a load-balancing cluster, use a URL of the form, **discovery**:

(`fabric:GroupName`), which automatically load balances the client across the available brokers in the cluster. For example, to connect a client to the `loadbal` cluster, you would use a URL like the following:

```
discovery:(fabric:loadbal)
```

For convenience, the `mq-create` command automatically generates a profile named `mq-client-GroupName`, which you can combine either with the `example-mq-consumer` profile or with the `example-mq-producer` profile to create a client of the load-balancing cluster.

For example, to create a consumer client of the `loadbal` group, you can deploy the `mq-client-loadbal` profile and the `example-mq-consumer` profile together in a child container, by entering the following command:

```
JBossFuse:karaf@root> container-create-child --profile mq-client-loadbal -
-profile example-mq-consumer root consumer
The following containers have been created successfully:
Container: consumer.
```

To create a producer client of the `loadbal` group, you can deploy the `mq-client-loadbal` profile and the `example-mq-producer` profile together in a child container, by entering the following command:

```
JBossFuse:karaf@root> container-create-child --profile mq-client-loadbal -
-profile example-mq-producer root producer
The following containers have been created successfully:
Container: producer.
```

To verify that the clients are functioning correctly, you can connect to one of them and check the log. For example, to check the log of the consumer client:

```
JBossFuse:karaf@root> container-connect consumer

JBossFuse:admin@consumer> log:display
2014-01-16 14:31:41,776 | INFO | Thread-42 | ConsumerThread
| org.fusesource.mq.ConsumerThread 54 | 110 - org.jboss.amq.mq-client -
6.1.0.redhat-312 | Received test message: 982
2014-01-16 14:31:41,777 | INFO | Thread-42 | ConsumerThread
| org.fusesource.mq.ConsumerThread 54 | 110 - org.jboss.amq.mq-client -
6.1.0.redhat-312 | Received test message: 983
```

7.3.2. Master-Slave Cluster


Overview

In the master-slave pattern, multiple peer brokers provide the same service and all compete to be the master. Only *one* master can exist at a given time, while the rest remain on standby as slaves. If the master stops, the remaining brokers (slaves) compete to become the new master. If the broker containers are deployed across different machines or data centres, the result is a highly available broker.

For example, [Figure 7.2, “Master-Slave Cluster”](#) shows a master-slave cluster with the group name, `masterslave`, and three brokers that compete with each other to register as the broker, `hq-broker`.

A broker becomes the master by acquiring a lock (where the lock implementation is provided by the underlying ZooKeeper registry). The other two brokers that fail to acquire the lock remain as slaves (but they continue trying to acquire the lock, at regular time intervals).

Figure 7.2. Master-Slave Cluster



Create brokers in a master-slave cluster

The basic rules for creating a master-slave cluster of brokers are as follows:

- Choose a group name for the master-slave cluster.
- Each broker in the cluster registers with the chosen group.
- Each broker must be identified by the *same* virtual broker name.
- Normally, each broker is deployed in a separate container.

For example, consider the cluster shown in [Figure 7.2, “Master-Slave Cluster”](#). The group name is `masterslave` and the cluster consists of three broker instances, each with the *same* broker name: `hq-broker`. You can create this cluster by entering a single `fabric:mq-create` command, as follows:

```
JBossFuse:karaf@root> mq-create --create-container broker --replicas 3 --
group masterslave hq-broker
```

Alternatively, if you have already created three containers, `broker1`, `broker2` and `broker3` (possibly running on separate machines), you can deploy a cluster of three brokers to the containers by entering the following command:

```
JBossFuse:karaf@root> mq-create --assign-container broker1,broker2,broker3
--group masterslave hq-broker
```

The first broker that starts becomes the master, while the others are slaves. When you stop the master, one of the slaves will take over and clients will reconnect. If brokers are persistent, you need to ensure that they all use the same store—for details of how to configure this, see [the section called “Configuring persistent data”](#).

Configure clients of a master-slave cluster

To connect a client to a master-slave cluster, use a URL of the form, `discovery:(fabric:GroupName)`, which automatically connects the client to the current master server. For example, to connect a client to the `masterslave` cluster, you would use a URL like the following:

```
discovery:(fabric:masterslave)
```

You can use the automatically generated client profile, `mq-client-masterslave`, to create sample clients. For example, to create an example consumer client in its own container, enter the following console command:

```
JBossFuse:karaf@root> container-create-child --profile mq-client-
masterslave --profile example-mq-consumer root consumer
The following containers have been created successfully:
Container: consumer.
```

And to create an example producer client in its own container, enter the following console command:

```
JBossFuse:karaf@root> container-create-child --profile mq-client-
masterslave --profile example-mq-producer root producer
The following containers have been created successfully:
Container: producer.
```

Locking mechanism

One benefit of this kind of master-slave architecture is that it does not depend on shared storage for locking, so it can be used even with non-persistent brokers. The broker group uses ZooKeeper to manage a shared distributed lock that controls ownership of the master status.

Re-using containers for multiple clusters

Fabric supports re-using the same containers for multiple master-slave clusters, which is a convenient way to economize on hardware resources. For example, given the three containers, `broker1`, `broker2`, and `broker3`, already running the `hq-broker` cluster, it is possible to reuse the *same* containers for another highly available broker cluster, `web-broker`. You can assign the `web-broker` profile to the existing containers with the following command:

```
mq-create --assign-container broker1,broker2,broker3 web-broker
```

This command assigns the new `web-broker` profile to the same containers already running `hq-broker`. Fabric automatically prevents two masters from running on the same container, so the master for `hq-broker` will run on a different container from the master for `web-broker`. This arrangement makes optimal use of the available resources.

Configuring persistent data

When you run a master-slave configuration with persistent brokers, it is important to specify where your store is located, because you need to be able to access it from multiple hosts. To support this scenario, the `fabric:mq-create` command enables you to specify the location of the data directory, as follows:

```
mq-create --assign-container broker1 --data /var/activemq/hq-broker hq-
broker
```

The preceding command creates the `hq-broker` virtual broker, which uses the `/var/activemq/hq-broker` directory for the data (and store) location. You can then mount some shared storage to this path and share the storage amongst the brokers in the master-slave cluster.

7.3.3. Broker Networks

Overview

It is possible to combine broker clusters with broker networks, giving you a hybrid broker network that combines the benefits of broker clusters (for example, high availability) with the benefits of broker networks (managing the flow of messages between different geographical sites).

Broker networks

A *broker network* in JBoss A-MQ is a form of federation where brokers are linked together using *network connectors*. This can be used as a way of forwarding messages between different geographical locations. Messages can be forwarded either *statically* (where specified categories of messages are always forwarded to a specific broker), or *dynamically* (where messages are forwarded only in response to a client that connects to a broker and subscribes to particular queues or topics).

For more details, see ["Using Networks of Brokers"](#) from the JBoss A-MQ library.


Creating network connectors

In the context of Fabric, network connectors can be created by passing the `--networks` option to the `fabric:mq-create` command.

Example broker network

Consider the scenario shown in [Figure 7.3, "Broker Network with Master-Slave Clusters"](#).

Figure 7.3. Broker Network with Master-Slave Clusters



The figure shows two master-slave clusters:

- The first cluster has the group name, `us-west`, and provides high-availability with a master-slave cluster of two brokers, `us-west1` and `us-west2`.
- The second cluster has the group name, `us-east`, and provides high-availability with a master-slave cluster of two brokers, `us-east1` and `us-east2`.

Network connectors link the master brokers between each of the geographical locations (there are, in fact, two network connectors in this topology: from west to east and from east to west).

To create the pair of master-slave brokers for the `us-east` group (consisting of the two containers `us-east1` and `us-east2`), you would log on to a root container running in the US East location and enter a command like the following:


```
mq-create --group us-east --networks us-west --networks-username User --
networks-password Pass --create-container us-east us-east
```

Where the `--networks` option specifies the name of the broker group you want to connect to, and the *User* and *Pass* are the credentials required to log on to the `us-west` broker cluster. By default, the `fabric:mq-create` command creates a master/slave pair of brokers.

And to create the pair of master-slave brokers for the `us-west` group (consisting of the two containers `us-west1` and `us-west2`), you would log on to a root container running in the US West location and enter a command like the following:

```
mq-create --group us-west --networks us-east --networks-username User --
networks-password Pass --create-container us-west us-west
```

Where *User* and *Pass* are the credentials required to log on to the `us-east` broker cluster.



NOTE

In a real scenario, you would probably first create the containers on separate machines and then assign brokers to the containers, using the `--assign-container` option in place of `--create-container`.

Connecting to the example broker network

At the US East location, any clients that need to connect to the broker network should use the following client URL:

```
discovery:(fabric:us-east)
```

And at the US West location, any clients that need to connect to the broker network should use the following client URL:

```
discovery:(fabric:us-west)
```

Any messages that need to be propagated between locations, from US East to US West (or from US West to US East), are transmitted over the broker network through one of the network connectors.

7.4. BROKER CONFIGURATION

Overview

The examples presented so far have demonstrated how to create brokers with default configuration settings. In practice, you will usually need to customize the broker configurations and this can be done by editing the properties of the corresponding Fabric profiles.

Setting OSGi Config Admin properties

Many of the broker configuration settings can be altered by editing OSGi Config Admin properties (which are organized into collections identified by a *persistent ID* or PID). For example, consider the `broker1` profile created by entering the following `fabric:mq-create` command:

```
fabric:mq-create --create-container broker --replicas 1 --networks us-west
```

brokerx

The preceding command creates the new profile, `mq-broker-default.brokerx`, and assigns this profile to the newly created `broker1` container.



NOTE

The new profile gets the name `mq-broker-Group.BrokerName` by default. If you want the profile to have the same name as the broker (which was the default in JBoss A-MQ version 6.0), you can specify the profile name explicitly using the `--profile` option.

You can inspect the details of the `mq-broker-default.brokerx` profile using the `fabric:profile-display` command, as follows:

```
JBossFuse:karaf@root> profile-display mq-broker-default.brokerx
  Profile id: broker1
  Version   : 1.0
  Parents   : mq-base
  Associated Containers :

  Container settings
  -----

  Configuration details
  -----
  PID: org.fusesource.mq.fabric.server-brokerx
  standby.pool default
  connectors openwire
  broker-name broker1
  data /opt/fuse-fabric/data/broker1
  config zk:/fabric/configs/versions/1.0/profiles/mq-base/broker.xml
  group default
  network us-west
```

Associated with the `org.fusesource.mq.fabric.server-brokerx` PID are a variety of property settings, such as `network` and `group`. You can now add more properties to this PID to customize the broker configuration.

Setting network connector properties

You can specify additional configuration for network connectors, where the property names have the form `network.NetworkPropName`. For example, to add the setting, `network.bridgeTempDestinations=false`, to the PID for `brokerx`, enter the following console command:

```
profile-edit --pid org.fusesource.mq.fabric.server-
brokerx/network.bridgeTempDestinations=false brokerx
```

The deployed broker dynamically detects the change to this property and updates the network connector on the fly.

Network connector properties by reflection

Fabric uses reflection to set network connector properties. That is, any PID property of the form `network.OptionName` can be used to set the corresponding `OptionName` property on the `org.apache.activemq.network.NetworkBridgeConfiguration` class. In particular, this implies you can set any of the following `network.OptionName` properties:

Property	Default	Description
<code>name</code>	<code>bridge</code>	Name of the network - for more than one network connector between the same two brokers, use different names
<code>userName</code>	<code>None</code>	Username for logging on to the remote broker port, if authentication is enabled.
<code>password</code>	<code>None</code>	Password for logging on to the remote broker port, if authentication is enabled.
<code>dynamicOnly</code>	<code>false</code>	If <code>true</code> , only activate a networked durable subscription when a corresponding durable subscription reactivates, by default they are activated on start-up.
<code>dispatchAsync</code>	<code>true</code>	Determines how the network bridge sends messages to the local broker. If <code>true</code> , the network bridge sends messages asynchronously.
<code>decreaseNetworkConsumerPriority</code>	<code>false</code>	If <code>true</code> , starting at priority <code>-5</code> , decrease the priority for dispatching to a network Queue consumer the further away it is (in network hops) from the producer. If <code>false</code> , all network consumers use same default priority (that is, <code>0</code>) as local consumers.
<code>consumerPriorityBase</code>	<code>-5</code>	Sets the starting priority for consumers. This base value will be decremented by the length of the broker path when <code>decreaseNetworkConsumerPriority</code> is set.

Property	Default	Description
networkTTL	1	The number of brokers in the network that messages and subscriptions can pass through (sets both messageTTL and consumerTTL)
messageTTL	1	The number of brokers in the network that messages can pass through.
consumerTTL	1	The number of brokers in the network that subscriptions can pass through (keep to 1 in a mesh).
conduitSubscriptions	true	Multiple consumers subscribing to the same destination are treated as one consumer by the network.
duplex	false	If true , a network connection is used both to produce <i>and</i> to consume messages. This is useful for hub and spoke scenarios, when the hub is behind a firewall, and so on.
prefetchSize	1000	Sets the prefetch size on the network connector's consumer. It must be greater than 0 , because network consumers do not poll for messages

Property	Default	Description
suppressDuplicateQueueSubscriptions	false	If true , duplicate subscriptions in the network that arise from network intermediaries are suppressed. For example, consider brokers A , B , and C , networked using multicast discovery. A consumer on A gives rise to a networked consumer on B and C . In addition, C networks to B (based on the network consumer from A) and B networks to C . When true , the network bridges between C and B (being duplicates of their existing network subscriptions to A) will be suppressed. Reducing the routing choices in this way provides determinism when producers or consumers migrate across the network as the potential for dead routes (stuck messages) are eliminated. The networkTTL value needs to match or exceed the broker count to require this intervention.
suppressDuplicateTopicSubscriptions	true	If true , duplicate network topic subscriptions (in a cyclic network) are suppressed.

Property	Default	Description
bridgeTempDestinations	true	<p>Whether to broadcast advisory messages for temporary destinations created in the network of brokers. Temporary destinations are typically created for request-reply messages. Broadcasting the information about temp destinations is turned on by default, so that consumers of a request-reply message can be connected to another broker in the network and still send back the reply on the temporary destination specified in the JMSReplyTo header. In an application scenario where most or all of the messages use the request-reply pattern, this generates additional traffic on the broker network, because every message typically sets a unique JMSReplyTo address (which causes a new temp destination to be created and broadcasted with an advisory message in the network of brokers).</p> <p>If you disable this feature, this network traffic can be reduced, but in this case the producers and consumers of a request-reply message need to be connected to the same broker. Remote consumers (that is, connected through another broker in your network) will not be able to send the reply message, but instead will raise a temp destination does not exist exception.</p>
alwaysSyncSend	false	<p>If true, non-persistent messages are sent to the remote broker using request/reply semantics instead of oneway message semantics. This setting affects both persistent and non-persistent messages the same way.</p>
staticBridge	false	<p>If true, the broker does not respond dynamically to new consumers. It uses only staticallyIncludedDestinations to create demand subscriptions.</p>

Property	Default	Description
<code>useCompression</code>	<code>false</code>	Compresses the message body when sending it over the network.
<code>advisoryForFailedForward</code>	<code>false</code>	If <code>true</code> , send an advisory message when the broker fails to forward the message to the temporary destination across the bridge.
<code>useBrokerNamesAsIdSeed</code>	<code>true</code>	Add the broker name as a prefix to connections and consumers created by the network bridge. It helps with visibility.
<code>gcDestinationViews</code>	<code>true</code>	If <code>true</code> , remove any MBeans for destinations that have not been used for a while.
<code>gcSweepTime</code>	<code>60000</code>	The period of inactivity in milliseconds, after which we remove MBeans.
<code>checkDuplicateMessagesOnDuplex</code>	<code>false</code>	If <code>true</code> , check for duplicates on the duplex connection.

Broker configuration file

Another important aspect of broker configuration is the ActiveMQ broker configuration file, which is specified as a Spring XML file. In the context of Fabric, this file is stored as a resource in the ZooKeeper registry in the `mq-base` profile. That is, in the ZooKeeper registry, the `broker.xml` file is stored in the following location:

```
/fabric/configs/versions/1.0/profiles/mq-base/broker.xml
```

This file has the following default contents:

```
<beans
  xmlns="http://www.springframework.org/schema/beans"
  xmlns:amq="http://activemq.apache.org/schema/core"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd
http://activemq.apache.org/schema/core
http://activemq.apache.org/schema/core/activemq-core.xsd">

  <!-- Allows us to use system properties and fabric as variables in
this configuration file -->
  <bean
class="org.springframework.beans.factory.config.PropertyPlaceholderConfigu
```

```

rer">
    <property name="properties">
        <bean
class="org.fusesource.mq.fabric.ConfigurationProperties"/>
    </property>
</bean>

    <broker xmlns="http://activemq.apache.org/schema/core"
brokerName="{broker-name}" dataDirectory="{data}" start="false">

        <destinationPolicy>
            <policyMap>
                <policyEntries>
                    <policyEntry topic="" producerFlowControl="true"
memoryLimit="1mb">
                        <pendingSubscriberPolicy>
                            <vmCursor />
                        </pendingSubscriberPolicy>
                    </policyEntry>
                    <policyEntry queue="" producerFlowControl="true"
memoryLimit="1mb">
                        </policyEntry>
                    </policyEntries>
                </policyMap>
            </destinationPolicy>

            <managementContext>
                <managementContext createConnector="false"/>
            </managementContext>

            <persistenceAdapter>
                <kahaDB directory="{data}/kahadb"/>
            </persistenceAdapter>

            <transportConnectors>
                <transportConnector name="openwire"
uri="tcp://0.0.0.0:0"/>
            </transportConnectors>
        </broker>

    </beans>

```

Note that some of the PID properties from the profile are substituted into this template (for example, **broker - name** and **data**) and it's important that you reuse them properly. The easiest way to edit this configuration is to use the Fuse Management Console (see ["Management Console User Guide"](#)) or the built-in profile text editor (see [Appendix A, Editing Profiles with the Built-In Text Editor](#)).

Additional broker configuration files

If you like, you can create additional broker configuration files in the **mq - base** profile, for example:

```
/fabric/configs/versions/1.0/profiles/mq-base/mybroker.xml
```

You can then use this custom **mybroker.xml** configuration by invoking the **fabric:mq-create** command with the **- -config** option, as follows:


```
fabric:mq-create --config mybroker.xml brokerx
```

The `--config` option assumes that the configuration file is stored in the current version of the `mq-base` profile, so you need to specify only the file name (that is, the full ZooKeeper path is not required).

PART II. FABRIC IN PRODUCTION

Abstract

Deepen your understanding and understand the principles of running Fabric in a production environment.

CHAPTER 8. FABRIC ENSEMBLE AND REGISTRY

Abstract

The Fabric ensemble and registry is a critical part of the Fabric infrastructure. In a production environment, it is particularly important to understand the correct approach to creating and maintaining a Fabric ensemble.

8.1. FABRIC REGISTRY

Overview

Fuse Fabric uses [Apache ZooKeeper](#) (a highly reliable distributed coordination service) as its registry for storing cluster configuration and node registration.

ZooKeeper is designed with consistency and high availability in mind, while protecting against network splits, using the concept of a server quorum. For example, you might run five ZooKeeper servers and, so long as you have a quorum (three or more servers available), the ZooKeeper cluster is reliable and not in a network split.

Registry structure

The structure of the registry is a tree-like structure, similar to a filesystem. Each node of the tree (a *znode*) can hold data and can have children.

For example, the following shows an outline of the registry structure:

```

fabric
|
+----registry (runtime registry)
|   |
|   +----containers
|   |   |
|   |   +----root
|   |
|   +----configs (configuration registry)
|   |   |
|   |   +----versions
|   |   |   |
|   |   |   +----1.0
|   |   |   |   |
|   |   |   |   +----profiles
|   |   |   |   |   |
|   |   |   |   |   +----default
|   |   |
|   |   +----containers

```

Parts of the registry

Conceptually, the Fabric registry consists of two main parts:

- *Configuration Registry*—the logical configuration of your fabric, which typically contains no physical machine information. It contains details of the applications to be deployed and their dependencies.
- *Runtime Registry*—contains details of how many machines are actually running, their physical location, and what services they are implementing.

Making the registry highly available

With a single container hosting the registry, high availability is not supported. In order to have a highly available Fabric registry, you need to replicate the registry on multiple containers (on different physical hosts). The common term used to describe a group of servers that replicate the Fabric registry is an *ensemble*.

8.2. ADMINISTERING A FABRIC ENSEMBLE

Recommendations for an ensemble in production

To assure high availability of the Fabric registry in a production environment, it is recommended that you observe the following guidelines for a Fabric ensemble:

- Deploy a minimum of five Fabric servers in production (if one server is taken down for maintenance, one other server can fail, and the Fabric registry will still be available).
- Fabric servers should be deployed on separate host machines.
- Each Fabric server should only have a Fabric registry agent deployed inside it. No other profiles should be deployed in it.
- The size of the ensemble should be fixed at the outset, and not changed later (if you subsequently add or remove containers from the ensemble, the ZooKeeper IP ports would be re-assigned).

Creating an ensemble

A Fabric ensemble is created in two stages, as follows:

1. Create an initial ensemble, consisting of one Fabric server.
2. Expand the ensemble, by adding an even number of containers.

Creating an initial ensemble

An initial ensemble is usually created by invoking the `fabric:create` console command (which converts the current container into a Fabric server, which is a sole member of the newly created ensemble). Alternatively, when creating a new container with the `fabric:container-create-ssh` or `fabric:container-create-cloud` commands, you can pass the `--ensemble-server` option.

For details of how to create an initial ensemble using the `fabric:create` command, see [Chapter 3, Creating a New Fabric](#).

Expanding the ensemble

Once you have an initial ensemble, consisting of one Fabric server, you can expand the ensemble by invoking the `fabric:ensemble-add` command. To expand the ensemble, perform the following steps:

1. Create some new managed containers in the current fabric, which you can then add to the ensemble. Use the default profile for these new containers. For a production environment, it is recommended that you create at least four new managed containers (must be an even number), each running on their own host.
2. While logged on to a container in the fabric, use the `fabric:ensemble-add` command to add the managed containers to the ensemble. For example, given the four managed containers, `container1`, `container2`, `container3`, and `container4`, you would enter the following command:

```
fabric:ensemble-add container1 container2 container3 container4
```



NOTE

Because the total number of containers in an ensemble must always be odd, you must specify an *even* number of containers to the `fabric:ensemble-add` command.

3. To check that the ensemble has been successfully created, invoke the `fabric:container-list` command.



IMPORTANT

Do not attempt to expand (or shrink) a Fabric ensemble in a production environment. When you add containers to (or remove containers from) an ensemble, the ZooKeeper IP ports are all re-assigned, which typically causes the containers in the fabric to lose connectivity with the ensemble.

Taking a Fabric server down for maintenance

If you need to perform any maintenance on the host where a Fabric server is running, you can do this while maintaining availability of the Fabric registry, so long as a quorum (more than half) of the Fabric servers are still running. To stop a Fabric server, simply invoke the `fabric:container-stop` command, specifying the name of the Fabric server.

In general, it is recommended to have at least five Fabric servers in the ensemble. Three is not an adequate number. For example, with three servers in the ensemble consider what happens when you take a Fabric server down for maintenance. The two remaining Fabric servers form a quorum, but there is now no tolerance for failure. If one of the remaining Fabric servers fails, the whole fabric fails. In order to maintain high availability during maintenance, it is therefore essential to have at least five Fabric servers in the ensemble.

CHAPTER 9. FABRIC AGENTS

Abstract

The Fabric agent, which is responsible for provisioning on each container instance, is a key component of the Fabric infrastructure. When it comes to troubleshooting the behavior of a Fabric container, it is valuable to have an understanding of what the Fabric agent does and how it works.

9.1. INTRODUCTION

Fabric agent

The *Fabric agent* is the part of Fabric that is responsible for applying profiles to containers. The agent can run in any container and its role is to retrieve profile information from the registry and apply them to the container.

To be specific, the Fabric agent performs the following actions:

1. Retrieves the profiles and versions assigned to the container on which it is running.
2. Reconfigures the container.
3. Calculates what needs to be installed, removed or updated on the container.
4. Performs the requisite install, remove, and update actions.

Agent modules

In reality, the Fabric agent is composed of the following two modules:

[fabric-configadmin]

The Fabric configuration admin bridge. Translates the registry information into configuration information.

[fabric-agent]

The deployment agent. Reads the translated configuration and provisions the container accordingly.

Often, the term, *agent*, refers just to the deployment agent ([fabric-agent] module), but here we discuss both of the agent modules and describe the role of each in some detail.

9.2. THE CONFIGURATION ADMIN BRIDGE

Overview

The configuration admin bridge is responsible for bridging between the ZooKeeper registry and the OSGi Configuration Admin service. After the bridge connects to the ZooKeeper registry, it discovers what version is assigned to the container, retrieves the appropriate versions of the profiles assigned to the container, translates the profiles into configuration data, and applies the profile data to the container.

Information in a profile

Profiles can contain two distinct kinds of information:

- *Configuration information*—which includes:
 - System configuration
 - OSGi configuration
- *Provisioning information*—which includes lists of:
 - Bundles
 - Karaf features
 - FABs

Actions performed

The configuration admin bridge reads all of the relevant profiles and creates an OSGi configuration to represent them. The provisioning and system information are then stored under the `io.fabric8.agent` PID (in the context of the OSGi Configuration Admin service, a PID is a named collection of property settings).

If an assigned profile belongs to a hierarchy (profile inheritance) or if multiple profiles are assigned to the container, the configuration admin bridge takes this into account, resolving any overlapping configuration settings to produce an overlay view of the profiles. There is only one `io.fabric8.agent` PID, even when there are multiple assigned profiles.

The output from the configuration admin bridge is just a set of key-value pairs stored under the `io.fabric8.agent` PID.

Configuration updates

The configuration admin bridge watches the Fabric registry for changes, so that any updates to the container's assigned profiles are tracked and immediately applied to the local container's OSGi configuration.

9.3. THE DEPLOYMENT AGENT

Actions performed

The *deployment agent* listens for local configuration changes on the `io.fabric8.agent` PID. Any change in that configuration will trigger the deployment agent.

When the deployment agent is triggered, it performs the following actions:

1. The deployment agent reads the whole `io.fabric8.agent` PID and calculates what bundles are to be installed in the container.
2. If the profiles assigned to the container specify any Karaf features or FABs, the deployment agent translates them into a list of bundles, so that the agent obtains a complete list of bundles to install.

3. The deployment agent compares the list of bundles to install with the list of bundles currently installed, in order to identify:
 - Bundles to uninstall,
 - Bundles to install,
 - Bundles to update.
4. The deployment agent then performs the bundle uninstalling, installing, and updating in the container.

Downloading artifacts

The deployment agent is capable of downloading artifacts from two different types of maven repository:

- Registered Fabric Maven proxies
- Configured Maven repositories (any Maven repository configured in the profile overlay).

Priority is always given to the Fabric Maven proxies. If more than one Maven proxy is registered in the fabric, the proxies are used in order, from the oldest to the newest.

If the target artifact is not found in the Maven proxies, the configured Maven repositories are used instead. The list of repositories is determined by the `org.ops4j.pax.url.mvn.repositories` property of the `io.fabric8.agent` PID.

To change the list of repositories for a specific profile, you can simply change the `org.ops4j.pax.url.mvn.repositories` property using the `fabric:profile-edit` command:

```
fabric:profile-edit --pid
io.fabric8.agent/org.ops4j.pax.url.mvn.repositories=http://repositorymanag
er.mylocalnetwork.net default
```

It is recommended that you specify this configuration in one profile only, and have the rest of profiles inherit from it. The `default` profile, which is the ancestor of all of the standard profiles, is the ideal place for this.

Container restarts

In most cases, when a container is provisioned by the provisioning agent, the container is kept alive and no restart is needed. A restart becomes necessary, however, whenever the following changes are made:

- Changes to the OSGi framework;
- Changes to the OSGi framework configuration.

The normal case is where the container stays alive during provisioning, because it is rarely necessary to make changes to the underlying OSGi framework. If a container does need to restart, the restart is performed automatically by the deployment agent and, after the restart, the container reconnects to the cluster without any manual intervention.

Monitoring the provisioning status

Throughout the whole process of deploying and provisioning, the deployment agent stores the provisioning status in the runtime registry, so that it is available to the whole cluster. The user can check the provisioning status at any time using the `fabric:container-list` command.

```
JBossFuse:karaf@root> fabric:container-list

[id]                [version] [alive] [profiles]
[provision status]
root*                1.0       true   fabric, fabric-
ensemble-0000-1 success
mq1                  1.0       true   mq
success
mq2                  1.0       true   mq
downloading
billing-broker       1.0       true   billing
success
  admin-console      1.0       true   web, admin-console
success
```

To monitor the provisioning status in real time, you can pass the `fabric:container-list` command as an argument to the `shell:watch` command, as follows:

```
shell:watch fabric:container-list
```

Resolution and startup ordering

To figure out what bundles need to be installed and what bundles need to be removed, the Fabric agent uses the OSGi Bundle Repository (OBR) resolver. The OBR resolver makes sure that all requirements are met (usually package requirements, but potentially also service requirements). To discover a bundle's requirements, the OBR reads the following bundle headers:

Import - Package

For each package listed here, the OBR resolver searches for a bundle that declares the package in a corresponding `Export - Package` header.

Import - Service

For each service listed here, the OBR resolver searches for a bundle that declares the service in a corresponding `Export - Service` header.

If you are using Blueprint configuration files, it is especially important to be aware of the need to add an `Export - Service` header to bundles that implement services. Blueprint configuration files with mandatory references to services will automatically be packaged with the `Import - Service` bundle header (assuming that you use the `maven-bundle-plugin`). If the bundle that exports the service does not explicitly specify an `Export - Service` header, resolution will fail. To fix this error, either the exporter bundle must add an `Export - Service` declaration, or the importer bundle must remove the `Import - Service` directive.

If resolution is successful, the Fabric agent will start the bundles. Even though you should try to avoid having requirements in the startup order of your bundles, the Fabric agent will attempt to start the bundles based on their expressed requirements and capabilities. This will not solve all issues, especially in cases where asynchronous service registration is involved. The best way to deal with this kind of issues is to use OSGi services.

CHAPTER 10. ALLOCATING PORTS

Abstract

You can use the port service to take care of allocating ports for your services, where the port service allocates ports in such a way as to avoid port clashes.

10.1. THE PORT SERVICE

What is the port service?

The port service is designed to address the problem of clashing IP port values, which frequently arises in a production environment. The following kinds of problem commonly arise:

- *Ports clashing with third-party services*—a server machine in a production environment often has multiple services deployed on it, with a wide range of IP ports in use. In this environment, there is a relatively large risk that a Fabric container could clash with existing IP ports.
- *Ports clashing with other Fabric containers*—when multiple Fabric containers are deployed on the same host, it is necessary to configure their standard services with different IP ports. Setting the IP ports manually would be a considerable nuisance (and error prone).
- *Ports clashing within a container*—a port clash can also occur within a single container, if multiple services are competing for the same ports (for example, multiple routes binding to the same ports). Because Fabric containers are highly dynamic, we need to be able to prevent port clashes in this case, and ports must be allocated and de-allocated as services come and go.

The port service addresses this problem by taking over the process of allocating ports. A service that uses the port service can specify a range of ports that it is willing to use, and the port service takes care of allocating a port that does not clash with any of the existing services.

Benefits of the port service

The port service offers the following benefits at run time:

- *Avoiding port clashes for standard container services*
- *Avoiding port clashes for custom services*

Avoiding port clashes for standard container services

When you start up multiple containers on the same host, the port service ensures that the containers automatically choose different IP ports for their standard services, thus avoiding port clashes between containers. You get this benefit for free: the standard container services are already configured to use the port service.

Avoiding port clashes for custom services

You can also use the port service for your own applications, enabling your custom services to avoid port clashes. This requires you to configure your custom services, as appropriate, to use the port service.

Using the port service in your own applications

To use the port service in your own application, proceed as follows:

1. Use the OSGi Config Admin service to define a key, whose value is a port range. Use the following syntax to define a key:

```
KeyID = ${port:MinValue,MaxValue}
```

The preceding syntax defines the key, *KeyID*, where *MinValue* specifies the minimum value of the IP port, and *MaxValue* specifies the maximum value of the IP port. You can create this key using the standard Karaf commands for editing persistent IDs (PIDs) and their keys (using the `fabric:profile-edit` command with the `--pid` option in a Fabric container).

For example, if you are logged into a Fabric container, you can see that the `default` profile defines the key, `org.osgi.service.http.port`, which specifies the container's Jetty port, as follows:

```
FuseFabric:karaf@root> fabric:profile-display default
...
PID: org.ops4j.pax.web
org.ops4j.pax.web.config.checksum ${checksum:profile:jetty.xml}
org.ops4j.pax.web.config.url profile:jetty.xml
javax.servlet.context.tempdir ${karaf.data}/pax-web-jsp
org.osgi.service.http.port ${port:8181,8282}
```

2. In your application's XML configuration (either Spring XML or Blueprint XML), replace the literal port value in the service's address by a property placeholder—for example, `${org.osgi.service.http.port}`—which substitutes the value of the key defined in step 1.

For a complete example of how to configure the property placeholder, see [Section 10.2, “Using the Port Service”](#).

How the port service allocates a port

Given a service with a port range (for example, `${port:9090,9190}`) running on a specific target host, when you start up the service for the *first time*, the port service allocates a port as follows:

1. Determines which ports in the range are already in use on the target host (whether local or remote), by actually trying to bind to the ports.
2. Checks the registered ports in the ZooKeeper registry for *all* of the containers deployed on the target host (even if the containers are currently not running).
3. Allocates the first free port, within the specified range, that does not clash with any of the ports discovered in steps 1 and 2.

How allocated ports are stored

Allocated ports are stored permanently in the ZooKeeper registry, under the following registry node:

```
/fabric/registry/ports/
```

Each key value, *KeyID*, is filed under its corresponding persistent ID, *PID*, and container name, *ContainerName*, as follows:

```
/fabric/registry/ports/containers/ContainerName/PID/KeyID
```

For example, given the child container, **Child1**, the key for the child container's Jetty port would be stored in the following ZooKeeper node:

```
/fabric/registry/ports/containers/Child1/org.ops4j.pax.web/org.osgi.service.http.port
```

Keys used by the standard container services

Some of keys used by standard container services are as follows:

```
/fabric/registry/ports/containers/ContainerName/org.apache.karaf.shell/sshPort  
/fabric/registry/ports/containers/ContainerName/org.ops4j.pax.web/org.osgi.service.http.port  
/fabric/registry/ports/containers/ContainerName/org.apache.karaf.management/rmiServerPort  
/fabric/registry/ports/containers/ContainerName/org.apache.karaf.management/rmiRegistryPort
```

Behavior upon stopping and restarting a container

When you stop a container, the ports used by that container are stored in the ZooKeeper registry and continue to be reserved for that container by the port service. Subsequently, when you restart the container, Fabric reads the port values stored in ZooKeeper and restarts the container's services using the stored values. This behavior has the following consequences:

- The ports used by the container's services remain constant (after the initial allocation has occurred). You can advertise the ports to clients and be confident that the ports will remain valid over the long term.
- If, while the container is stopped, another service binds to one of the container's ports, there is a port clash when the container restarts, and the affected service fails to start (but at least we can guarantee that Fabric will not cause such a clash, because Fabric deliberately avoids re-using allocated container ports).

Deallocating ports

When you destroy a container (by invoking the `fabric:container-delete` command), Fabric deallocates all of the ports assigned to that container, so that they become available for use again by services in other containers. In other words, when the *ContainerName* container is deleted, all of the key entries under `/fabric/registry/ports/containers/ContainerName` are deleted from the ZooKeeper registry.

10.2. USING THE PORT SERVICE

Overview

This section explains how to use the port service in your own applications, taking the `example-camel-cxf` profile as an example. There are two basic steps to configuring the port service in your application:

- *At development time*—using the property placeholder service, replace a service's fixed port number by a key.
- *At deployment time*—using the OSGi Config Admin service, specify the key value as a port range. For example, you can specify the key value as a PID property setting in a Fabric profile.

It is possible to configure the property placeholder in Blueprint XML, or in Java (using the relevant OSGi API).



NOTE

The property placeholder syntax in Spring XML is deprecated (it belongs to the deprecated Spring-DM component).

Demonstration code

This example is based on the `example-camel-cxf` profile. The source code for the example is taken from the `fabric-camel-cxf` example on [Github](#), which is available from the following URL:

```
https://github.com/fabric8io/fabric8/tree/1.x/fabric/fabric-examples/fabric-camel-cxf
```

Property placeholder in XML configuration

The following Spring XML configuration shows the definition of an endpoint for the greeter Web service (taken from the file, `src/main/resources/OSGI-INF/blueprint/cxf.xml`, in the `fabric-camel-cxf` demonstration):

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:cxf="http://camel.apache.org/schema/blueprint/cxf"
  xmlns:cm="http://aries.apache.org/blueprint/xmlns/blueprint-
cm/v1.1.0"
  xsi:schemaLocation="
  http://www.osgi.org/xmlns/blueprint/v1.0.0
  http://www.osgi.org/xmlns/blueprint/v1.0.0/blueprint.xsd
  http://camel.apache.org/schema/blueprint/cxf
  http://camel.apache.org/schema/blueprint/cxf/camel-cxf.xsd">

  <cm:property-placeholder id="placeholder"
    persistent-
id="io.fabric8.examples.camel.cxf"
    update-strategy="reload">
    <cm:default-properties>
      <cm:property name="greeterPort" value="9090"/>
    </cm:default-properties>
  </cm:property-placeholder>

  <cxf:cxfEndpoint id="greeterEndpoint"
    address="http://localhost:${greeterPort}/greeter"
```

```

                serviceClass="io.fabric8.examples.camelcxf.Greeter">
    <cxf:features>
      <bean class="io.fabric8.cxf.endpoint.ManagedApiFeature"/>
    </cxf:features>
  </cxf:cxfEndpoint>
  ...
</blueprint>

```

The CXF endpoint (which binds to a Camel route) is defined by the `cxf:cxfEndpoint` element. In the `address` attribute, the port number is specified by substituting the `greeterPort` key, `${greeterPort}`. The property placeholder mechanism is configured by the `cm:property-placeholder` element, which specifies that the `greeterPort` property belongs to the `io.fabric8.examples.camel.cxf` PID. The property placeholder mechanism is integrated with the OSGi Config Admin service, which allows you to override the port number at deployment time.

Specifying a port range using OSGi Config Admin

At deployment time, you can override the default port number of the greeter Web service. In this particular example, where the deployment is described by the `example-camel-cxf` profile, the port number is integrated with the port service and specified as a port range.

Because the port range is defined at deployment time, it is *not* specified in the example source code, but is instead specified in the `example-camel-cxf` Fabric profile. You can see the configured port range by entering the following console command:

```
JBossFuse:karaf@root> fabric:profile-display example-camel-cxf
```

In the output of this command, you should see the following configuration setting for the `io.fabric8.examples.camel.cxf` persistent ID:

```

...
Configuration details
-----
PID: io.fabric8.examples.camel.cxf
    greeterPort ${port:9090,9190}
...

```

The preceding output shows that the `greeterPort` key is set to `${port:9090,9190}`.

Modifying the port range

If you want to modify the port range configured in the `example-camel-cxf` profile, you can do so using the `fabric:profile-edit` console command. For example, to change the value of `greeterPort` to the range, `${port:7070,7170}`, you would enter the following console command:

```
JBossFuse:karaf@root> fabric:profile-edit
  --pid io.fabric8.examples.camel.cxf/greeterPort=\${port:7070,7170}
  example-camel-cxf
```

Where the `$` sign and the curly braces, `{ }`, must be escaped by the backslash character, `\`, as shown. Alternatively, if you prefer to edit the port range using the built-in text editor, you can enter the following console command instead:

■

```
JBossFuse:karaf@root> fabric:profile-edit --pid  
io.fabric8.examples.camel.cxf example-camel-cxf
```

CHAPTER 11. GATEWAY

Abstract

The Fabric Gateway provides a TCP and HTTP/S gateway for discovery, load balancing and failover of services running in a fabric. The Fabric Gateway enables you to use standard HTTP URLs to access Web applications or Web services running in a fabric. In JBoss A-MQ, messaging clients can discover and connect to brokers over any supported messaging protocol (OpenWire, STOMP, MQTT, AMQP or WebSockets), letting the gateway handle the connection management to the real services running inside the fabric.



WARNING

The Fabric Gateway is a technical preview only, and is *not* supported in JBoss A-MQ 6.1.

11.1. GATEWAY ARCHITECTURE

Deployment strategies

There are two main deployment strategies for a gateway:

- Run the gateway on each machine that needs to discover services and communicate with it through `localhost`. In this case, you do not need to hard code any host names in your messaging or Web clients and the connection to the gateway on `localhost` is nice and fast.
- Run the gateway on one or more known hosts using DNS or VIP load balancing (mapping host names to machines). In this case, you can use a fixed host name for all your services

How the gateway works

The gateway monitors and detects any changes in the ZooKeeper registry for all Web applications, Web services, servlets and message brokers. For all of the registered services, the gateway applies mapping rules to figure out how to expose those services through TCP or HTTP.

The ZooKeeper registry is automatically populated by Fabric when you deploy Web archives (WARs) or CXF based Web services.

11.2. RUNNING THE GATEWAY

Deploy a gateway profile

To run the gateway, simply deploy one (or more) of the predefined profiles to a Fabric container. The following gateway profiles are provided:

`gateway - mq`

Profile for a messaging gateway (for accessing Apache ActiveMQ brokers in the fabric).

gateway-http

Profile for a HTTP gateway (for Web applications or Web services).

11.3. CONFIGURING THE GATEWAY

Configuring with the Management Console

To configure the gateway using the Management Console UI, navigate to the **Profiles** page then click on the **Configuration** tab, then select either the **Fabric8 HTTP Gateway** or the **Fabric8 MQ Gateway** to configure its settings.

HTTP mapping rules

When using the HTTP gateway, it is a common requirement to map different versions of Web applications or Web services to different URI paths on the gateway. You can perform very flexible mappings using [URI templates](#).

The default behavior is to expose all Web applications and Web services at the context path they are running in the target server. For example, if you deploy the `example-quickstarts-rest` profile, that uses a URI like `/cxf/crm/customerservice/customers/123` on whatever host and port it is deployed on. Hence, by default, it is visible on the gateway at <http://localhost:9000/cxf/crm/customerservice/customers/123>. For this example, the URI template is:

```
{contextPath}/
```

Which means take the context path (in the above case, `/cxf/crm`) and append `/`, giving `/cxf/crm/`. Any request within that path is then passed to an instance of the CXF `crm` service.

Selecting part of the ZooKeeper registry

The mapping rules for the MQ gateway and the HTTP gateway are tied to particular regions of the ZooKeeper registry. If you specify a ZooKeeper path for a mapping rule, any services registered under that path become associated with that rule.

For example, in the case of messaging, you could associate a messaging gateway with all message brokers worldwide. Alternatively, you could provide continent-specific, country-specific or region-specific gateways, just by specifying different ZooKeeper paths for each gateway configuration. For regional messaging clusters, use different ZooKeeper folders for geographically distinct broker clusters.

With HTTP then REST APIs, SOAP Web Services, servlets and web applications all live in different parts of the ZooKeeper registry. From the Management Console UI, you can browse the contents of the registry in the **Runtime | Registry** section of the console (in the **Fabric** view).

Here are the common ZooKeeper paths:

ZooKeeper Path	Description
----------------	-------------

<code>/fabric/registry/clusters/apis/rest</code>	REST based web services
<code>/fabric/registry/clusters/apis/ws</code>	SOAP based web services
<code>/fabric/registry/clusters/servlets</code>	Servlets (registered usually individually via the OSGI APIs)
<code>/fabric/registry/clusters/webapps</code>	Web Applications (i.e. WARs)

Segregating URI paths

You might want to segregate servlets, Web services, or Web applications into different URI spaces.

For example, if you want all Web services to be available under `/api/` and Web applications to be available under `/app/`, update the URI templates as follows:

For the Web services mapping rule:

```
ZooKeeperPath: /fabric/registry/clusters/apis
URI template: /api{contextPath}/
```

For the Web applications mapping rule:

```
ZooKeeperPath: /fabric/registry/clusters/webapps
URI template: /app{contextPath}/
```

If you want to split RESTful APIs and SOAP web services into different URI paths, replace the preceding mapping rule with the following rules:

```
ZooKeeperPath: /fabric/registry/clusters/apis/rest
URI template: /rest{contextPath}/

ZooKeeperPath: /fabric/registry/clusters/apis/ws
URI template: /ws{contextPath}/
```

11.4. VERSIONING

Explicit URIs

You might want to expose all available versions of each Web service and Web application at a different URI. For example, consider the case where you change your URI template to the following:

```
/version/{version}{contextPath}/
```

If you have **1.0** and **1.1** versions of a profile that packages Web services or Web applications, you can now access the different versions using version-specific URIs. For example, if you are running version **1.0** and version **1.1** implementations of the `example-quickstarts-rest` profile, you can access either one through the following URIs:

- Version 1.0 through
<http://localhost:9000/version/1.0/cxf/crm/customerservice/customers/123>
- Version 1.1 through
<http://localhost:9000/version/1.1/cxf/crm/customerservice/customers/123>

Both versions are available to the gateway, provided you include the version information in the URI.

Rolling upgrades

Another approach to dealing with versions of Web services and Web applications is to expose only a *single* version at a time of each Web service or Web application in a single gateway. This is the default configuration.

For example, if you deploy a **1.0** version of the `gateway-ht tp` profile and run a few services, you will see all **1.0** versions of them. If you run some **1.1** versions of these services, the gateway will not see them. If you now do a rolling upgrade of your gateway to version **1.1**, it will switch to showing only the **1.1** versions of the services.

Alternatively, you can specify the exact *profile* version to use, on the mapping configuration screen.

Another approach you can use with Web applications is to specify the maven coordinates and maven version of a web application in the ZooKeeper path.

11.5. URI TEMPLATE EXPRESSIONS

Variables

The following table shows the variables you can use in a URI template expression:

Expression	Description
<code>{bundleName}</code>	The name of the bundle that registers the Web service, servlet or application. This variable is currently <i>not</i> supported for Web services, but works for Web applications and servlets in an OSGi container.
<code>{bundleVersion}</code>	The version of the bundle that registers the Web service, servlet or application. This variable is currently <i>not</i> supported for Web services, but works for Web applications and servlets in an OSGi container.
<code>{container}</code>	The container ID of the container where the Web service or Web application is deployed.
<code>{contextPath}</code>	The context path (the part of the URL after the host and port) of the Web service or Web application implementation.

{servicePath}	The relative path within ZooKeeper that a service is registered. This is usually is made up of, for web services as the service name and version. For web applications its often the maven coordinates
{version}	The profile version of the Web service or Web application.

CHAPTER 12. SECURING FABRIC CONTAINERS

Abstract

By default, fabric containers uses text-based username/password authentication. Setting up a more robust access control system involves creating and deploying a new JAAS realm to the containers in the fabric.

DEFAULT AUTHENTICATION SYSTEM

By default, Fabric uses a simple text-based authentication system (implemented by the JAAS login module, `io.fabric8.jaas.ZookeeperLoginModule`). This system allows you to define user accounts and assign passwords and roles to the users. Out of the box, the user credentials are stored in the Fabric registry, unencrypted.

MANAGING USERS

You can manage users in the default authentication system using the `jaas:*` family of console commands. First of all you need to attach the `jaas:*` commands to the `ZookeeperLoginModule` login module, as follows:

```
JBossFuse:karaf@root> jaas:realms
Index Realm          Module Class
  1 karaf
org.apache.karaf.jaas.modules.properties.PropertiesLoginModule
  2 karaf
org.apache.karaf.jaas.modules.publickey.PublickeyLoginModule
  3 karaf          io.fabric8.jaas.ZookeeperLoginModule
JBossFuse:karaf@root> jaas:manage --index 3
```

Which attaches the `jaas:*` commands to the `ZookeeperLoginModule` login module. You can then add users and roles, using the `jaas:useradd` and `jaas:roleadd` commands. Finally, when you are finished editing the user data, you must commit the changes by entering the `jaas:update` command, as follows:


```
JBossFuse:karaf@root> jaas:update
```

Alternatively, you can abort the pending changes by entering `jaas:cancel`.

OBFUSCATING STORED PASSWORDS

By default, the JAAS `ZookeeperLoginModule` stores passwords in plain text. You can provide additional protection to passwords by storing them in an obfuscated format. This can be done by adding the appropriate configuration properties to the `io.fabric8.jaas` PID and ensuring that they are applied to *all* of the containers in the fabric.

For more details, see [section "Encrypting Stored Passwords" in "Security Guide"](#).

**NOTE**

Although message digest algorithms are not easy to crack, they are not invulnerable to attack (for example, see the [Wikipedia article on cryptographic hash functions](#)). Always use file permissions to protect files containing passwords, in addition to using password encryption.

ENABLING LDAP AUTHENTICATION

Fabric supports LDAP authentication (implemented by the Apache Karaf `LDAPLoginModule`), which you can enable by adding the requisite configuration to the default profile.

For details of how to enable LDAP authentication in a fabric, see [chapter "LDAP Authentication Tutorial" in "Security Guide"](#).

CHAPTER 13. CONFIGURING A FABRIC'S MAVEN PROXY

Abstract

The Fabric Ensemble creates a Maven proxy to access the repositories from which artifacts are distributed to the fabric's containers. You can modify the default settings to use a different set of repositories or the make an internal repository accessible.

OVERVIEW

The Fabric Ensemble creates a Maven proxy to facilitate access to the artifacts required by the containers in the fabric. Each Fabric Server deployed in the fabric runs an instance of a Maven proxy. The ensemble aggregates all of the proxies so that it appears to the Fabric Agents as a single Maven proxy.

The Fabric Agents use the fabric's Maven proxy to access the known repositories. This ensures that all of the containers use the same set of repositories and artifacts.



NOTE

Advanced users can configure each Fabric Server to act as a proxy for a different set of repositories. However, this is not a recommended set up.



NOTE

Red Hat JBoss Fuse Tooling provides tooling for uploading bundles using the Maven proxy. You can also add the fabric's Maven Proxy to a POM file so that bundles can be distributed to the ensemble as part of an automated build process.

DEFAULT REPOSITORIES

By default a fabric's Maven proxy is configured to be a proxy for the following Maven repositories:

- Maven Central (<http://repo1.maven.org/maven2>)
- Fuse Releases (<http://repo.fusesource.com/nexus/content/repositories/releases>)
- Fuse Early Access (<http://repo.fusesource.com/nexus/content/groups/ea>)
- SpringSource (<http://repository.springsource.com/maven/bundles/release>,
<http://repository.springsource.com/maven/bundles/external>)
- Scala Tools (<http://scala-tools.org/repo-releases>)
- User's Local (`~/.m2/repository`)

CHANGING THE REPOSITORIES

To change the repositories the ensemble proxies:

1. Create a new profile version.

From the command console this is done using the `fabric:version-create` command. See [section "Description" in "Console Reference" section "Description" in "Console Reference"](#) for more information.

2. Change the `org.ops4j.pax.url.mvn.repositories` property in the `io.fabric8.agent` PID of the `default` profile. [Example 13.1, "Configuring the Maven Proxy URL"](#) shows the console command for editing this property.

Example 13.1. Configuring the Maven Proxy URL

```
JBossFuse:karaf@root> fabric:profile-edit -p
io.fabric8.agent/org.ops4j.pax.url.mvn.repositories= \
    http://repo1.maven.org/maven2, \
    http://repo.fusesource.com/nexus/content/repositories/releases, \
    http://repo.fusesource.com/nexus/content/groups/ea, \
    http://repository.springsource.com/maven/bundles/release, \
    http://repository.springsource.com/maven/bundles/external, \
    http://scala-tools.org/repo-releases default
```



NOTE

The `io.fabric8.agent` PID is refined in all of the fabric profiles. Setting the proxy URL, the `org.ops4j.pax.url.mvn.repositories` property, in the `default` profile ensures that all of the other fabric profiles share the same Maven proxy setting.



IMPORTANT

The fabric profile's `io.fabric8.maven` PID, which ultimately controls the Maven proxy, imports its value from the `default` profile's `io.fabric8.agent` PID. You should **not** change the settings of the `io.fabric8.maven` PID.

3. Roll the changes out the fabric by upgrading the containers to the new profile version.



IMPORTANT

You cannot test this configuration change out on a few containers to validate it. The change **must** be made to the entire fabric or it will result in conflicts.

CHAPTER 14. OFFLINE REPOSITORIES

Abstract

Its quite common a common requirement to need offline repositories: either as a local cache of remote Maven repositories, or in cases where production machines do not have access to the Internet.

14.1. OFFLINE REPOSITORY FOR A PROFILE

Download into a specified directory

To download all the bundles and features of a given profile, *ProfileName*, enter the following console command:

```
fabric:profile-download --profile ProfileName /tmp/myrepo
```

This command downloads all the bundles, FABs and features for the default version of the given profile into the */tmp/myrepo* directory.

Download into the system folder

If you omit the path, the `fabric:profile-download` command installs the files to the `system` folder inside the current Fuse container (thereby populating the local maven repository for the container). For example:

```
fabric:profile-download --profile ProfileName
```

14.2. OFFLINE REPOSITORY FOR A VERSION

Download the current version

To download all the bundles and features for all the profiles in the default version, enter the following console command:

```
fabric:profile-download /tmp/myrepo
```

Download a specific version

You can specify the version to download using the `--version` option, as follows:

```
fabric:profile-download --version 1.0 /tmp/myrepo
```

If you omit the path the `fabric:profile-download` command installs the files into the `system` folder inside the current Fuse container (thereby populating the local maven repository for the container).

14.3. OFFLINE REPOSITORY FOR A MAVEN PROJECT

Download repository for Maven project

If you have a Maven project and you need to create an offline repository for building this project and its runtime dependencies, you can use the [maven dependency plugin](#).

For example, from the top-level directory of a Maven project (such that the current directory has a `pom.xml` file), you should be able to run the following Maven command:

```
mvn org.apache.maven.plugins:maven-dependency-plugin:2.8:go-offline -  
Dmaven.repo.local=/tmp/cheese
```

Which downloads all the Maven dependencies and plug-ins required to build the project to the `/tmp/cheese` directory.

CHAPTER 15. PATCHING

15.1. PATCHING A CONTAINER IN A FABRIC

Abstract

In a fabric patches are applied to profiles and the patched version of the profile is applied to the container. The management console is the recommended tool for patching containers in a fabric. The `fabric` shell also has the commands needed to apply a patch and roll it out to running containers.

Overview

The bundles loaded by a container in a fabric are controlled by the container's Fabric Agent. The agent inspects the profiles applied to the container to determine what bundles to load, and the version of each bundle, and then loads the specified version of each bundle for the container.

A patch typically includes a new version of one or more bundles, so to apply the patch to a container in a fabric you need to update the profiles applied to it. This will cause the Fabric Agent to load the patched versions of the bundles.

The management console is the recommended tool for patching containers in a fabric. However, the command console's `fabric` shell also provides the commands needed to patch containers running in a fabric.

Procedure

Patching a container in a fabric involves:

1. Getting a patch file.
 - o Customer Support sends you a patch.
 - o Customer Support sends you a link to download a patch.
 - o You, or your organization, generate a patch file for an internally created application.
2. Uploading one or more patch files to the fabric's Maven repository.
3. Applying the patch(es) to a profile version.

This creates a new profile version that points to the new versions of the patched bundles and repositories.

4. Migrate one or two containers to the patched profile version to ensure that the patch does not introduce any new issues.
5. After you are certain that the patch works, migrate the remaining containers in the fabric to the patched version.

Using the management console

The management console is the easiest and most verbose method of patching containers in a fabric. Its `Patching` tab uploads patches to a fabric's Maven repository and applies the patch to a specified

profile version. You can then use the management console to roll the patch out to all of the containers in the fabric.


See [chapter "Patching a Fabric" in "Management Console User Guide"](#) for more information.

Using the command console

The Red Hat JBoss Fuse command console can also be used to patch containers running in a fabric. To patch a fabric container:

1. Create a new version, using the `fabric:version-create` command:

```
JBossFuse:karaf@root> fabric:version-create 1.1
Created version: 1.1 as copy of: 1.0
```



IMPORTANT

The version name must be a pure *numeric* string, such as **1.1**, **1.2**, **2.1**, or **2.2**. You cannot incorporate alphabetic characters in the version name (such as **1.0.patch**).

2. Apply the patch to the new version, using the `fabric:patch-apply` command. For example, to apply the `activemq.zip` patch file to version **1.1**:

```
JBossFuse:karaf@root> fabric:patch-apply --version 1.1
file:///patches/activemq.zip
```

3. Upgrade the container using the `fabric:container-upgrade` command, specifying which container you want to upgrade. For example, to upgrade the `root` container, enter the following command:

```
JBossFuse:karaf@root> fabric:container-upgrade 1.1 root
Upgraded container root from version 1.0 to 1.1
```

4. You can check that the new patch profile has been created using the `fabric:profile-list` command, as follows:

```
JBossFuse:karaf@root> fabric:profile-list --version 1.1 | grep patch
default                0                patch-
activemq-patch
patch-activemq-patch
```

Where we presume that the patch was applied to profile version 1.1.

TIP

If you want to avoid specifying the profile version (with `--version`) every time you invoke a profile command, you can change the default profile version using the `fabric:version-set-default Version` command.

You can also check whether specific JARs are included in the patch, for example:

```
JBossFuse:karaf@root> list | grep -i activemq
[ 131] [Active      ] [Created      ] [      ] [ 50] activemq-osgi
(5.9.0.redhat-61037X)
[ 139] [Active      ] [Created      ] [      ] [ 50] activemq-
karaf (5.9.0.redhat-61037X)
[ 207] [Active      ] [      ] [      ] [ 60] activemq-
camel (5.9.0.redhat-61037X)
```

APPENDIX A. EDITING PROFILES WITH THE BUILT-IN TEXT EDITOR

Abstract

When you have a lot of changes and additions to make to a profile's configuration, it is usually more convenient to do this interactively, using the built-in text editor for profiles. The editor can be accessed by entering the `profile-edit` command with no arguments except for the profile's name (and optionally, version); or adding the `--pid` option for editing OSGi PID properties; or adding the `--resource` option for editing general resources.

A.1. EDITING AGENT PROPERTIES

Overview

This section explains how to use the built-in text editor to modify a profile's *agent properties*, which are mainly used to define what bundles and features are deployed by the profile.

Open the agent properties resource

To start editing a profile's agent properties using the built-in text editor, enter the following console command:

```
JBossFuse:karaf@root> profile-edit Profile [Version]
```

Where *Profile* is the name of the profile to edit and you can optionally specify the profile version, *Version*, as well. The text editor opens in the console window, showing the current profile name and version in the top-left corner of the Window. The bottom row of the editor screen summarizes the available editing commands and you can use the arrow keys to move about the screen.

Specifying feature repository locations

To specify the location of a feature repository, add a line in the following format:

```
repository.ID=URL
```

Where *ID* is an arbitrary unique identifier and *URL* gives the location of a single feature repository (only one repository URL can be specified on a line).

Specifying deployed features

To specify a feature to deploy (which must be available from one of the specified feature repositories), add a line in the following format:

```
feature.ID=FeatureName
```

Where *ID* is an arbitrary unique identifier and *FeatureName* is the name of a feature.

Specifying deployed bundles

To specify a bundle to deploy, add a line in the following format:

```
bundle.ID=URL
```

Where **ID** is an arbitrary unique identifier and **URL** specifies the bundle's location.



NOTE

A bundle entry can be used in combination with a **blueprint:** (or **spring:**) URL handler to deploy a Blueprint XML resource (or a Spring XML resource) as an OSGi bundle.

Specifying bundle overrides

To specify a bundle override, add a line in the following format:

```
override.ID=URL
```

Where **ID** is an arbitrary unique identifier and **URL** specifies the bundle's location.



NOTE

A bundle override is used to override a bundle installed by a feature, replacing it with a different version of the bundle. For example, this functionality is used by the patching system to install a patched bundle in a container.

Specifying deployed FABs

To specify a Fuse Application Bundle (FAB) to deploy, add a line in the following format:

```
fab.ID=URL
```

Where **ID** is an arbitrary unique identifier and **URL** specifies the FAB's location.

Specifying etc/config.properties properties

To specify Java system properties that affect the Apache Karaf container (analogous to editing **etc/config.properties** in a standalone container), add a line in the following format:

```
config.Property=Value
```

Specifying etc/system.properties properties

To specify Java system properties that affect the bundles deployed in the container (analogous to editing **etc/system.properties** in a standalone container), add a line in the following format:

```
system.Property=Value
```

If the system property, **Property**, is already set at the JVM level (for example, through the **--jvm-opts** option to the **fabric:container-create** command), the preceding **fabric:profile-edit** command *will not override the JVM level setting*. To override a JVM level setting, set the system

property as follows:

```
system.karaf.override.Property=Value
```

Specifying libraries to add to Java runtime lib/

To specify a Java library to deploy (equivalent to adding a library to the `lib/` directory of the underlying Java runtime), add a line in the following format:

```
lib.ID=URL
```

Where **ID** is an arbitrary unique identifier and **URL** specifies the library's location.

Specifying libraries to add to Java runtime lib/ext/

To specify a Java extension library to deploy (equivalent to adding a library to the `lib/ext/` directory of the underlying Java runtime), add a line in the following format:

```
ext.ID=URL
```

Where **ID** is an arbitrary unique identifier and **URL** specifies the extension library's location.

Specifying libraries to add to Java runtime lib/endorsed/

To specify a Java endorsed library to deploy (equivalent to adding a library to the `lib/endorsed/` directory of the underlying Java runtime), add a line in the following format:

```
endorsed.ID=URL
```

Where **ID** is an arbitrary unique identifier and **URL** specifies the endorsed library's location.

Example

To open the `mq-client` profile's agent properties for editing, enter the following console command:

```
JBossFuse:karaf@root> profile-edit mq-client
```

The text editor starts up, and you should see the following screen in the console window:

```
Profile:mq-client 1.0
L:1 C:1
#
# Copyright (C) Red Hat, Inc.
# http://redhat.com
#
# Licensed under the Apache License, Version 2.0 (the "License");
# you may not use this file except in compliance with the License.
# You may obtain a copy of the License at
#
#    http://www.apache.org/licenses/LICENSE-2.0
#
```



```
# Unless required by applicable law or agreed to in writing, software
# distributed under the License is distributed on an "AS IS" BASIS,
# WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
# See the License for the specific language governing permissions and
# limitations under the License.
#

repository.activemq=mvn:org.apache.activemq/activemq-
karaf/${version:activemq}/xml/features
repository.karaf-
standard=mvn\ :org.apache.karaf.assemblies.features/standard/${version:kara
f}/
xml/features

      ^X Quit      ^S Save      ^Z Undo      ^R Redo      ^G Go To      ^F Find
^N Next      ^P Previous
```

Type `^X` to quit the text editor and get back to the console prompt.

A.2. EDITING OSGI CONFIG ADMIN PROPERTIES

Overview

This section explains how to use the built-in text editor to edit the property settings associated with a specific persistent ID.

Persistent ID

In the context of the OSGi Config Admin service, a *persistent ID* (PID) refers to and identifies a set of related properties. In particular, when defining PID property settings in a profile, the properties associated with the *PID* persistent ID are defined in the *PID.properties* resource.

Open the Config Admin properties resource

To start editing the properties associated with the *PID* persistent ID, enter the following console command:

```
JBossFuse:karaf@root> profile-edit --pid PID Profile [Version]
```



NOTE

It is also possible to edit PID properties by specifying `--resource PID.properties` in the `profile-edit` command, instead of using the `--pid PID` option.

Specifying OSGi config admin properties

The text editor opens, showing the contents of the specified profile's *PID.properties* resource (which is actually stored in the ZooKeeper registry). To edit the properties, add, modify, or delete lines of the following form:

```
Property=Value
```

Example

To edit the properties for the `io.fabric8.hadoop` PID in the `hadoop-base` profile, enter the following console command:

```
JBossFuse:karaf@root> profile-edit --resource io.fabric8.hadoop.properties
hadoop-base 1.0
```

The text editor starts up, and you should see the following screen in the console window:

```
Profile:hadoop-base 1.0
L:1 C:1
#
# Copyright (C) Red Hat, Inc.
# http://redhat.com
#
# Licensed under the Apache License, Version 2.0 (the "License");
# you may not use this file except in compliance with the License.
# You may obtain a copy of the License at
#
#     http://www.apache.org/licenses/LICENSE-2.0
#
# Unless required by applicable law or agreed to in writing, software
# distributed under the License is distributed on an "AS IS" BASIS,
# WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
# See the License for the specific language governing permissions and
# limitations under the License.
#

fs.default.name=hdfs\://localhost\:9000
dfs.replication=1
mapred.job.tracker=localhost\:9001
dfs.name.dir=${karaf.data}/hadoop/dfs/name
dfs.http.address=0.0.0.0\:9002
dfs.data.dir=${karaf.data}/hadoop/dfs/data
dfs.name.edits.dir=${karaf.data}/hadoop/dfs/name

      ^X Quit      ^S Save      ^Z Undo      ^R Redo      ^G Go To      ^F Find
^N Next      ^P Previous
```

You might notice that colon characters are escaped in this example (as in `\:`). Strictly speaking, it is only necessary to escape a colon if it appears as part of a property name (left hand side of the equals sign), but the `profile-edit` command automatically escapes all colons when it writes to a resource. When manually editing resources using the text editor, however, you do not need to escape colons in URLs appearing on the right hand side of the equals sign.

Type `^X` to quit the text editor and get back to the console prompt.

A.3. EDITING OTHER RESOURCES

Overview

In addition to agent properties and PID properties, the built-in text editor makes it possible for you edit *any* resource associated with a profile. This is particularly useful, if you need to store additional

configuration files in a profile. The extra configuration files can be stored as profile resources (which actually correspond to ZooKeeper nodes) and then can be accessed by your applications at run time.



NOTE

The ZooKeeper registry is designed to work with small nodes only. If you try to store a massive configuration file as a profile resource, it will severely degrade the performance of the Fuse Fabric registry.

Creating and editing an arbitrary resource

You can create and edit arbitrary profile resources using the following command syntax:

```
JBossFuse:karaf@root> profile-edit --resource Resource Profile [Version]
```

Where **Resource** is the name of the profile resource you want to edit. If **Resource** does not already exist, it will be created.

broker.xml example

For example, the **mq-base** profile has the **broker.xml** resource, which stores the contents of an Apache ActiveMQ broker configuration file. To edit the **broker.xml** resource, enter the following console command:

```
JBossFuse:karaf@root> profile-edit --resource broker.xml mq-base 1.0
```

The text editor starts up, and you should see the following screen in the console window:

```
Profile:mq-base 1.0
L:1 C:1
<!--
  Copyright (C) FuseSource, Inc.
  http://fusesource.com

  Licensed under the Apache License, Version 2.0 (the "License");
  you may not use this file except in compliance with the License.
  You may obtain a copy of the License at

      http://www.apache.org/licenses/LICENSE-2.0

  Unless required by applicable law or agreed to in writing, software
  distributed under the License is distributed on an "AS IS" BASIS,
  WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
  See the License for the specific language governing permissions and
  limitations under the License.
-->
<beans
  xmlns="http://www.springframework.org/schema/beans"
  xmlns:amq="http://activemq.apache.org/schema/core"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd
http://activemq.apache.org/schema/core
http://activemq.apache.org/schema/core/activemq-core.xsd">
```

```

    <!-- Allows us to use system properties and fabric as variables in
this configuration file -->
    <bean
class="org.springframework.beans.factory.config.PropertyPlaceholderConfigu
rer">
        <property name="properties">
            <bean
class="org.fusesource.mq.fabric.ConfigurationProperties"/>
        </property>
    </bean>
^X Quit      ^S Save      ^Z Undo      ^R Redo      ^G Go To      ^F Find
^N Next      ^P Previous

```

Any changes you make to this file will take effect whenever the broker restarts.

Type `^X` to quit the text editor and get back to the console prompt.

Referencing a profile resource

In order to use an arbitrary profile resource, you must be able to reference it. Because a profile resource is ultimately stored as a ZooKeeper node, you must reference it using a ZooKeeper URL. For example, the `broker.xml` resource from the previous example is stored under the following ZooKeeper location:

```
zk:/fabric/configs/versions/1.0/profiles/mq-base/broker.xml
```

In general, you can find version, *Version*, of the *Profile* profile's *Resource* resource at the following location:

```
zk:/fabric/configs/versions/Version/profiles/Profile/Resource
```

For example, the `mq` profile's `org.fusesource.mq.fabric.server-broker` PID defines the following properties, where the `config` property references the `broker.xml` resource:

```

connectors=openwire
config=zk\:/fabric/configs/versions/1.0/profiles/mq-base/broker.xml
group=default
standby.pool=default

```

A.4. PROFILE ATTRIBUTES

Overview

In addition to the resources described in the other sections, a profile also has certain attributes that affect its behavior. *You cannot edit profile attributes directly using the text editor.*

For completeness, this section describes what the profile attributes are and what console commands you can use to modify them.

parents attribute

The `parents` attribute is a list of one or more parent profiles. This attribute can be set using the

`profile-change-parents` console command. For example, to assign the parent profiles `camel` and `cx` to the `my-camel-cxf-profile` profile, you would enter the following console command:

```
JBossFuse:karaf@root> profile-change-parents --version 1.0 my-camel-cxf-  
profile camel cxf
```

abstract attribute

When a profile's `abstract` attribute is set to `true`, the profile cannot be directly deployed to a container. This is useful for profiles that are only intended to be the parents of other profiles—for example, `mq-base`. You can set the abstract attribute from the Management Console.

locked attribute

A locked profile cannot be changed or edited until it is unlocked. You can lock or unlock a profile from the Management Console.

hidden attribute

The `hidden` attribute is a flag that is typically set on profiles that Fuse Fabric creates automatically (for example, to customize the setup of a registry server). By default, hidden profiles are not shown when you run the `profile-list` command, but you can see them when you add the `--hidden` flag, as follows:

```
JBossFuse:karaf@root> profile-list --hidden  
...  
fabric 1 karaf  
fabric-ensemble-0000 0  
fabric-ensemble-0000-1 1 fabric-ensemble-0000  
fmc 0 default  
...
```

APPENDIX B. FABRIC URL HANDLERS

Abstract

The Fabric runtime provides a variety of URL handlers, which can be used in application code deployed in a Fabric-enabled container. These URLs are intended to be used in profile configuration files to locate configuration resources.

B.1. PROFILE URL HANDLER

The profile URL is used to access resources stored under the current profile (or parent profile). It has the following format:

```
profile:ResourceName
```

A key characteristic of the profile URL is that the location of a resource can change dynamically at run time, as follows:

1. The profile URL handler first tries to find the named resource, *ResourceName*, in the *current version* of the *current profile* (where the current version is a property of the container in which the profile is running).
2. If the specified resource is not found in the current profile, the profile URL tries to find the resource in the current version of the *parent profile*.

This behavior implies that whenever you change the version assigned to a container (for example, by invoking the `fabric:container-upgrade` or `fabric:container-rollback` console commands), the referenced resources are also, automatically, upgraded or rolled back.

B.2. ZK URL HANDLER

You can reference the contents of a ZooKeeper node using the zk URL. The URL can be specified either as an absolute node:

```
zk:/PathToNode
```

Or you can reference configuration properties from a specific container using the following syntax:

```
zk:ContainerName/Property
```

The preceding syntax is effectively a short cut to the following URL reference:

```
zk:/fabric/registry/containers/config/ContainerName/Property
```

B.3. BLUEPRINT URL HANDLER

The Blueprint URL handler enables you to deploy a Blueprint XML resource directly as an OSGi bundle, without needing to create any of the usual OSGi bundle packaging in advance. The `blueprint:` scheme can be prefixed to any of the usual location URL handlers (for example, `file:`, `http:`, `profile:`, `zk:`).

To use the Blueprint URL handler, create a **bundle** entry in the agent properties (equivalent to the **io.fabric8.agent** PID) in the following format:

```
bundle.ID=blueprint:LocationScheme:LocationOfBlueprintXML
```

For example, the **example-camel-hello** profile's agent properties include the following **bundle** entry:

```
bundle.camel-fabric=blueprint:profile:camel.xml
```



NOTE

The Blueprint URL handler has an important side effect. If the referenced Blueprint resource is changed at run time, the Blueprint URL handler detects this change and automatically reloads the resource. This means, for example, that if you edit a deployed Camel route in a Blueprint resource, the route automatically gets updated in real time.

B.4. SPRING URL HANDLER

The Spring URL handler enables you to deploy a Spring XML resource directly as an OSGi bundle, without needing to create any of the usual OSGi bundle packaging in advance. The **spring:** scheme can be prefixed to any of the usual location URL handlers (for example, **file:**, **http:**, **profile:**, **zk:**).

To use the Spring URL handler, create a **bundle** entry in the agent properties (equivalent to the **io.fabric8.agent** PID) in the following format:

```
bundle.ID=spring:LocationScheme:LocationOfBlueprintXML
```

For example, to load the Spring resource, **camel-spring.xml**, from the current profile, you could add the following entry to the profile's agent properties:

```
bundle.spring-resource=spring:profile:camel-spring.xml
```



NOTE

If the referenced Spring resource is changed at run time, the Spring URL handler detects this change and automatically reloads the resource.

APPENDIX C. PROFILE PROPERTY RESOLVERS

Abstract

When defining properties for a profile, you can use a variable substitution mechanism, which has the general syntax `${ResourceReference}`. This variable substitution mechanism can be used in *any* profile resource, including the agent properties, PID properties, and other resources—for example, the mq-base profile's `broker.xml` resource references the `${broker.name}` and `${data}` variables.

C.1. SUBSTITUTING SYSTEM PROPERTIES

Syntax

System properties can be substituted in a profile resource, using the following syntax:

```
${PropName}
```

Where *PropName* can be the name of any Java system property. In particular, Java system properties can be defined in the following locations:

- The `etc/system.properties` file, relative to the container's home directory.
- System property settings in the profile's agent properties.

Some of the more useful system properties defined in the `etc/system.properties` file are, as follows:

Table C.1. System Properties

System Property	Description
<code>\${karaf.home}</code>	The directory where Red Hat JBoss Fuse is installed.
<code>\${karaf.data}</code>	Location of the current container's data directory, which is usually <code>\${karaf.home}/data</code> for a main container or <code>\${karaf.home}/instances/InstanceName/data</code> for a child container.
<code>\${karaf.name}</code>	The name of the current container.

C.2. SUBSTITUTING ENVIRONMENT VARIABLES

Syntax

You can substitute the value of a system environment variable using the environment property resolver, which has the following syntax:

```
${env:VarName}
```


C.3. SUBSTITUTING CONTAINER ATTRIBUTES

Syntax

You can substitute the value of a container attribute using the container attribute property resolver, which has the following syntax:

```
${container:Attribute}
```

You can substitute any of the following container attributes:

Table C.2. Container Attributes

Attribute	Description
<code>\${container:resolver}</code>	The effective <i>resolver policy</i> for the current container. Possible values are: localip , localhostname , publicip , publichostname , manualip .
<code>\${container:ip}</code>	The effective IP address used by the current container, which has been selected by applying the resolver policy. This is the form of host address that is advertised to other containers and applications.
<code>\${container:localip}</code>	The numerical IP address of the current container, which is suitable for accessing the container on a LAN.
<code>\${container:localhostname}</code>	The host name of the current container, which is suitable for accessing the container on a LAN.
<code>\${container:publicip}</code>	The numerical IP address of the current container, which is suitable for accessing the container from a WAN (on the Internet).
<code>\${container:publichostname}</code>	The host name of the current container, which is suitable for accessing the container from a WAN (on the Internet).
<code>\${container:manualip}</code>	An IP address that is specified manually, by setting the value of the relevant node in the ZooKeeper registry.
<code>\${container:bindaddress}</code>	
<code>\${container:sshurl}</code>	The URL of the SSH service, which can be used to log on to the container console.
<code>\${container:jmxurl}</code>	The URL of the JMX service, which can be used to monitor the container.

Attribute	Description
<code>#{container:jolokiaurl}</code>	The URL of the Jolokia service, which is used by the Fuse Management Console to access the container.
<code>#{container:httpurl}</code>	The base URL of the container's default Jetty HTTP server.
<code>#{container:domains}</code>	List of JMX domains registered by the container.
<code>#{container:processid}</code>	Returns the process ID of the container process (on Linux-like and UNIX-like operating systems).
<code>#{container:openshift}</code>	A boolean flag that returns true , if the container is running on OpenShift; otherwise, false .
<code>#{container:blueprintstatus}</code>	The aggregate status of all the deployed Blueprint contexts. If all of the deployed contexts are ok, the status is ok ; if one or more deployed contexts have failed, the status is failed .
<code>#{container:springstatus}</code>	The aggregate status of all the deployed Spring contexts. If all of the deployed contexts are ok, the status is ok ; if one or more deployed contexts have failed, the status is failed .
<code>#{container:provisionstatus}</code>	Returns the container provision status.
<code>#{container:provisionexception}</code>	If the container provisioning has failed, this variable returns the provisioning exception.
<code>#{container:provisionlist}</code>	The list of provisioned artefacts in the container.
<code>#{container:geolocation}</code>	The geographic location of the container (which is obtained by making a Web request to a public service that gives the GPS coordinates of the container host).

C.4. SUBSTITUTING PID PROPERTIES

Syntax

The profile property resolver is used to access PID properties from the current profile (or parent profile). It has the following format:

```
#{profile:PID/Property}
```



NOTE

This should not be confused with the syntax of a profile URL, which is used to access general resource files (not PID properties) and which is not resolved immediately (in contrast to the profile property resolver, which substitutes the corresponding property value as soon as the configuration file is read).

Example using a profile property resolver

For example, the `fabric` profile's `io.fabric8.maven.properties` PID resource includes the following property setting:

```
remoteRepositories=${profile:io.fabric8.agent/org.ops4j.pax.url.mvn.repositories}
```

So that the `remoteRepositories` property is set to the value of the `org.ops4j.pax.url.mvn.repositories` agent property (`io.fabric8.agent` is the PID for the agent properties).

C.5. SUBSTITUTING ZOOKEEPER NODE CONTENTS

Syntax

You can substitute the contents of a ZooKeeper node using the `zk` property resolver. The property resolver can be specified either as an absolute node:

```
${zk:/PathToNode}
```

Or you can reference configuration properties from a specific container using the following syntax:

```
${zk:ContainerName/Property}
```

The preceding syntax is effectively a short cut to the following property resolver:

```
${zk:/fabric/registry/containers/config/ContainerName/Property}
```

Recursive variable substitution

It is also possible to use a variable within a variable (recursive substitution). For example, the `dosgi` profile's `io.fabric8.dosgi.properties` resource defines the following property:

```
exportedAddress=${zk:${karaf.name}/ip}
```

How to reference the current version of a resource

A potential problem arises with ZooKeeper property resolver if you need to reference a ZooKeeper node that has a version number embedded in it. For example, suppose you want to reference the `my-profile` profile's `my-resource` resource, which can be done using the following ZooKeeper URL:

```
${zk:/fabric/configs/versions/1.0/profiles/my-profile/my-resource}
```

Notice that the profile version number, **1.0**, is embedded in this path. But if you decide to upgrade this profile to version **1.1**, this means you must manually edit all occurrences of this ZooKeeper URL, changing the version number to **1.1** in order to reference the upgraded resource. To avoid this extra work, and to ensure that the resolver always references the current version of the resource, you can use the following trick which exploits recursive variable substitution:

```
${zk:/fabric/configs/versions/${zk:/fabric/configs/containers/${karaf.name}}/profiles/my-profile/my-resource}
```

This works because the `/fabric/configs/containers/${karaf.name}` ZooKeeper node contains the current profile version deployed in the container.

C.6. CHECKSUM PROPERTY RESOLVER

Syntax

The checksum property resolver can be used, if you want a resource to reload automatically at run time, whenever it is updated. The `checksum:` scheme can be prefixed to any of the usual location URL handlers (for example, `file:`, `http:`, `profile:`, `zk:`).

For example, the `default` profile defines the following checksum property in the `org.ops4j.pax.web` PID:

```
org.ops4j.pax.web.config.checksum=${checksum:profile\:jetty.xml}
```

C.7. PORT PROPERTY RESOLVER

Syntax

The port property resolver is used to access the *port service*, which can automatically allocate an IP port within a specified range. It has the following syntax:

```
${port:Min,Max}
```

Where *Min* and *Max* specify the minimum and maximum values of the allocated IP port.

APPENDIX D. TECHNOLOGY-SPECIFIC DISCOVERY MECHANISMS

Abstract

In addition to the core provisioning, configuration, and management features described above, Fabric also provides a number of technology specific extensions.

D.1. ACTIVEMQ ENDPOINT DISCOVERY

If you were to deploy ActiveMQ brokers into a collection of standalone containers, you would have to provide an explicit list of broker locations to any messaging clients. In the case of a cloud deployment, where servers are frequently shut down and restarted in different locations, this type of configuration is very difficult to maintain.

Deploying the ActiveMQ brokers on a fabric, however, is much easier to manage, because you can exploit the Fabric discovery mechanism, which enables location transparency. When they start up, the brokers register themselves with the Fabric registry. Clients can then consult the Fabric registry to discover brokers. Moreover, because the discovery mechanism automatically detects when brokers fail and when new brokers start, there is also support for real-time discovery and load balancing.

For example, in the context of a Fabric deployment, it is typically sufficient to configure clients with a discovery URL like the following:

```
discovery:(fabric:us-east)
```

Where `us-east` represents a cluster of brokers deployed in the `us-east` cloud.

D.2. CAMEL ENDPOINT DISCOVERY

When deploying Camel endpoints in a cloud, it is often useful to enable location transparency for some Camel endpoints. In a dynamic cloud environment, where server machines are frequently shut down and restarted in a different location, you might want to have the capability to reconnect routes across different hosts. This kind of location transparency and reconnection functionality is supported by two special URI prefixes, which integrate Camel endpoints with the Fabric discovery mechanism:

fabric:

The `fabric:` prefix provides location transparency and load-balancing functionality for Camel endpoints. A consumer endpoint (acting like a server port), defines its endpoint URI using the following syntax:

```
fabric:ClusterID:EndpointURI
```

So that the Camel endpoint URI, *EndpointURI*, is stored in the Fabric registry under the cluster ID, *ClusterID*. Producer endpoints (acting like clients) can then access the Camel endpoint by specifying just the *ClusterID*, as follows:

```
fabric:ClusterID
```

Fabric then resolves the cluster ID to the registered endpoint URI. If multiple endpoint URIs are registered under the same cluster ID, Fabric randomly selects one of the available URIs, thus providing load-balancing functionality.

master :

The `master :` prefix provides failover functionality for Camel endpoints. Two or more consumer endpoints (server ports) must register their URIs with the Fabric registry, using the following syntax:

```
master:ClusterID:EndpointURI
```

Producer endpoints (acting like clients) can then access the failover cluster by defining a URI with the following syntax:

```
master:ClusterID
```

Initially, Fabric chooses one of the consumer endpoints to be the master instance and always resolves the client URIs to this master location. If the master fails, however, Fabric will fail over to one of the other registered endpoint URIs and start resolving to that endpoint instead.

D.3. CXF ENDPOINT DISCOVERY

For connections between Web service clients and Web services *within* the cloud, you have the option of supporting location transparency and load balancing through the Fabric discovery mechanism. This mechanism is not supported for connections coming into the cloud, however, so that you cannot advertise a Web service externally using this discovery mechanism.

To enable Fabric discovery on the server side, you must install the requisite *Fabric load balancer* feature into the CXF bus object. When this feature is installed, all Web services created in this bus will automatically be registered in the Fabric registry.

On the client side, you must install the fabric load balancer feature directly in the client proxy instance. The Web service client will then look up the server's location in the Fabric registry, ignoring the address configured in the client endpoint (you can use a dummy value for the URL).

D.4. OSGI SERVICE DISCOVERY

The OSGi registry typically stores references to local services. However, it is also relatively easy to configure Fabric to store *remote* OSGi services in the Fabric registry, so that you can realize invisible remoting of OSGi services, with support for dynamic discovery, load balancing and fail over. Fabric follows the *OSGi Remote Services* specification: any service exported to the OSGi registry with the `service.exported.interfaces` property (having a value of `*` or an explicit list of interfaces to export) is automatically exported to the Fabric registry and becomes accessible to other remote containers in the same fabric.

For example, to expose an OSGi service in the fabric using Blueprint XML, define the OSGi service as follows:

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0">
  <bean id="theBean" class="foo.bar.Example" />
  <service ref="theBean" auto-export="interfaces">
    <service-properties>
```

```
    <entry key="service.exported.interfaces" value="*" />
  </service-properties>
</service>
</blueprint>
```

To import an OSGi service from the fabric, simply reference the OSGi service in the usual way (for example, using Blueprint XML or through the OSGi Java API).